Design of Vibrotactile Gradients for Guiding the Manual Motion of an RFID-Enabled Hand-Held Locator towards Sought Items in Shelf-Sized Spaces

Krispin Johan Leydon

Master of Science in Computer Science

Interaction Design Centre
Department of Computer Science &
Information Systems
University of Limerick
Limerick, Ireland

Supervisors: Niall Griffith, Mikael Fernström

This thesis is presented in fulfillment of the requirements for a Master of Science degree in Computer Science. The work is my own, and has not been submitted previously to this or any other university or educational institution. Others' work has been fully acknowledged and referenced throughout.

Submitted to the University of Limerick, July 2004

Abstract

Hand-held appliances for associating physical items with digital data (such as bar code readers) serve an increasingly diverse range of applications, however their interactive feedback remains remarkably homogenous: the ubiquitous beep & blink. This feedback may accurately reflect such appliances' underlying digital implementations, however it does not support activities specifically, nor does it do justice to the rich heritage of continuous haptic feedback that has characterized the use of hand-tools for millennia.

This thesis explores the design, development and evaluation of active haptic feedback for one such associative appliance: a hand-held locator for finding items in shelf-sized spaces. More specifically, the thesis examines how gradients in vibrotactile feedback can be designed and delivered so as to guide manual motion towards sought items.

Work began with conceptualizing the locator (through scenarios) and verifying that it could be implemented (through a technology exploration). Next, vibrotactile feedback for the appliance was simulated and demonstrated to identify key dimensions of the interaction. Finally, experiments were conducted to evaluate the impact that variations along key dimensions had on locating time.

The study found that continuous gradients in vibrotactile feedback supported locating more effectively than discrete feedback in both qualitative and quantitative terms. The performance advantages of one continuous gradient over another for locating were inconclusive, as were the relative advantages of different ratios between the size of a sought item and the maximum range of vibrotactile feedback. Fitts' law was preliminarily explored, and found not to describe the experimental data. The concept of locating through vibrotactile feedback was validated, and points of departure found for further experimentation, implementation and application.

Acknowledgements

There are many people to thank for their assistance with this research: Niall Griffith for his patience, constancy and level-headed advice; Mikael Fernström for his encouragement and support; Liam Bannon for introducing a world of perspectives on the human-technical relationship; Stephen Hurley for his insightful and enthusiastic criticism; Anne Murphy for so patiently and professionally taking care of all sorts of etc. in the background; the rest of the Interaction Design Centre for their support and humor; Noel Guinane and James Keane for their generous loan of laboratory equipment; Tom Sorensen for sharing his expertise in inductive power transfer; Sile O'Modhrain and Gunnar Jansson for commenting on the project's initial proposal and final experimental design. This research was made possible through grants from the Higher Education Authority of Ireland.

Table of Contents

1. **Introduction** 1
2. **Background** 2
   2.1. Virtual and Embodied Approaches to Human Computer Interaction 2
   2.2. Physical-Digital Association 10
   2.3. “Focal Gradients” and Orienting Toward Tools for Mediated Display 26
   2.4. Haptic Display 29
   2.5. Fitts’ Law for Targeting 33
   2.6. Environmental Concerns and Computation 34
3. **Path to Research Topic** 39
   3.1. Early Decisions 39
   3.2. Research Topic (Rough Cut) 47
4. **Preliminary Explorations** 49
   4.1. Scenarios 49
   4.2. Technology Exploration 52
   4.3. Simulation of Functionality 56
   4.4. Demonstration 68
   4.5. Parameterization and Idealizations 75
   4.6. Tight Formulation of Research Questions 78
5. **Experimentation** 80
   5.1. Experiment Design 80
   5.2. The Experiments 92
   5.3. Results & Discussion 94
6. **Conclusions** 128
   6.1. Experimental Conclusions 128
   6.2. Practical Implications for the Design of Feedback for a Hand-Held Locator 129
   6.3. Limitations 131
7. **Future Work** 133
   7.1. Experimentation 133
   7.2. Implementation 134

7.3. Application 134
8. **References** 136

Appendix A: RFID Reader Schematic (Reference Design) A-1
Appendix B: Modified RFID Reader Design Notes B-1
Appendix C: Microcontroller Code for Modified RFID Reader C-1
Appendix D: Lessons from Implementing an RFID Reader D-1
Appendix E: Apparatus and Procedure for Inductor Measurement and Resonator Tuning. E-1
Appendix F: Legality of a Hand-Held RFID locator with Range Exceeding a Meter. F-1
Appendix G: Form-Giving: Sketches & Doodles G-1
Appendix H: Verifying the Linearity of the Vibrotactile Transducer H-1
Appendix I: Latency Measurement I-1
Appendix J: Script for 1st Experiment J-1
Appendix K: Questionnaire to be Completed After the First Experiment K-1
Appendix L: Script for 2nd Experiment L-1
Appendix M: Questionnaire to be Completed After the Second Experiment ... M-1
Appendix N: Additional Questionnaire Responses N-1
Appendix O: Declaration of Informed Consent for Participation in Time-Trial Experiments O-1
Appendix P: Relative Frequency of Mapping Function Orderings, Data Split by Mapping-Range P-1

Equations, Tables and Figures

Equation 1: Fitts' law..34

Table 1: Summary of Experimental Conditions…………………………………….91
Table 2: Participants of the first and second experiments......................................102
Table 3: Participants ranked by over-all mean find-time 105

Figure 1: High-level block diagram of an RFID system..15
Figure 2: Frequency response for high and low-Q resonators. 18
Figure 3: A parallel resonator 18
Figure 4: An RFID tag. 19
Figure 5: A switched parallel resonator 20
Figure 6: A transistor-switched parallel resonator. 20
Figure 7: Interrogator distance as a function of magnetic field strength. 23
Figure 8: Relative orientation of interrogator antenna and tag inductor...................24
Figure 9: "SweptRF" tag tracking systems 25
Figure 10: A system with negative feedback 27
Figure 11: The seamless transision enabled by an ordinary magnifying glass. 28
Figure 12: The PHANTOM® Desktop™ 32
Figure 13: High-level block diagram of an RFID reader...53
Figure 14: Implemented RFID tag reader 54
Figure 15: Revised transmitter subsection. 54
Figure 16: High-level block diagram of the simulatory apparatus 57
Figure 17: VBW32 transducer used for vibrotactile display and custom amplifier..58
Figure 18: Polhemus Fastrak sensor system 58
Figure 19: Vibrotactile mouse during development...60
Figure 20: Simulation GUI..62
Figure 21: Scavenged shapes 63
Figure 22: Hand-held locator's final form 64
Figure 23: Sample mapping functions. 66
Figure 24: Posture for locating...67

Figure 25: Demo participant locating books through vibrotactile feedback............ 70
Figure 26: Ratio of sought item size to mapping function range............................. 72
Figure 27: Experimental objects and parameters.. 81
Figure 28: Experimental environment & apparatus... 82
Figure 29: Actual experimental apparatus and environment.................................. 84
Figure 30: Mapping functions chosen for experimentation..................................... 86
Figure 31: Participants in the midst of time trial experiments................................ 94
Figure 32: Reconstructed movement patterns of the hand-held locator.................... 97
Figure 33: Writable "mix" CDs – an example of homogeneous packaging........... 100
Figure 34: Percentages of targeting hits and misses/cancels for all trials.............. 104
Figure 35: The average normalized skew of participants' find-time distributions.. 107
Figure 36: Relative frequency of the mapping function orderings........................ 108
Figure 37: Pair-wise comparisons between the find-time distributions of each mapping function for all participants at a confidence interval of 95%........... 110
Figure 38: Pair-wise comparisons between the find-time distributions of each mapping function for all participants at a confidence interval of 80%........... 111
Figure 39: Relative frequency of the mapping function orderings, for two different values of mapping-range.. 113
Figure 40: Relative frequency of two-way mapping function orderings............... 115
Figure 41: Relative frequency of two-way mapping function orderings at 95% and 80% confidence intervals. 115
Figure 42: Sample plots of find-time as a function of the ratio of target-size to mapping-range.. 117
Figure 43: Curves used to visually classify participants' mean find-time as a function of the ratio between target-size and mapping-range...................................... 117
Figure 44: Expected and examined (hypothetical) relationship between find-time and the ratio of target-size to mapping-range 118
Figure 45: Sample plots of find-time as a function of mapping-range.................. 120
Figure 46: Curves used to visually classify participants' mean find-time as a function of mapping-range.. 121
Figure 47: Sample plots of find-time as a function of target-home distance.......... 123

1. Introduction

Warm,
Cold, colder, cold...
Warm - getting warmer.
Hot. Very hot - You've found it!

This reconstructed exchange from the popular children's game “find the spoon” (known by various names in various places) illustrates how a gradient in feedback - in this case a “hot” center surrounded by an increasingly “cool” periphery - can assist a seeker to locate an object of search. In a variety of search and targeting activities (such as aiming a flashlight, tuning an analog radio and positioning a magnifying glass) we experience continuous gradients in intensity, gradients that guide our alignment with the phenomena we seek to observe.

But where does “cold” end? Where should “hot” begin? How might guiding gradients be designed, and how might their delivery to our senses be automated so as to support locating things in the physical world?

This thesis chronicles the design of guiding gradients for a hand-held locator for finding items in shelf-sized spaces. It provides a personal account of one full turn of a design cycle - from concept to evaluation - that is pragmatic with respect to implementation, and speculative with regard to application.

The structure of the thesis is as follows: First, the Background chapter introduces prerequisite technical, historical and conceptual foundations. Next, the Path to Research Topic chapter draws upon these foundations to orient the reader towards the design space to be explored. The Preliminary Explorations chapter discusses a rapid exploration of this space conducted in order to choose a focus for the more detailed work discussed subsequently, in the chapter on Experimentation. The Conclusions chapter summarizes the lessons learned along with their implications and limitations, and the Future Work chapter recommends next steps that might be taken.

2. Background

This chapter provides a backdrop of concerns, ideas and agendas against which the activities and decisions of later chapters will make sense. First, it discusses over-arching perspectives on the design of interactive systems. Next it introduces some of the technical concerns involved in coordinating physical things with digital representations, and discusses in detail the strengths, weaknesses and potential of one particular associative technology. The subsequent section addresses a feature that is common to traditional tools we look, hear and feel through, and discusses this feature's implications for *digital* tools for looking, hearing and feeling through. The following two sections discuss the means and motivations for "haptic" human-computer interaction and a mathematical model for describing targeting activities, respectively. Finally, the chapter discusses some of computation's more ominous impacts on built and natural environments. While these topics may seem disparate, each has a bearing on the decisions made in the coming chapters, and all must be understood for the overall trajectory of the research to make sense.

2.1. Virtual and Embodied Approaches to Human Computer Interaction

While our physical bodies, things and spaces may matter to us, the physical world goes almost wholly unregistered by computational processes. This disconnection often forces us to choose – between digital capabilities on the one hand, and physical sensibilities on the other. Pierre Wellner, a researcher at Xerox' EuroParc, describes this problem as it applies to documents: "each world [physical and digital] has advantages and constraints that lead us to chose one or the other for particular tasks. Unfortunately, choosing to interact with a document in one world means forgoing the advantages of the other" (Wellner, 1993) p87. Hiroshi Ishii, of the MIT Media Lab's Tangible Media Group, describes the divide in terms of citizenship:

> We live between two realms: our physical environment and cyberspace. Despite our dual citizenship, the absence of seamless couplings between these parallel existences leaves a

> great divide between the worlds of bits and atoms. At the present, we are torn between these parallel but disjoint spaces (Ishii and Ullmer, 1997) p234.

This "two worlds problem", as it has become known, has proven to be a knotty one philosophically as well as technically, and research in human computer interfaces can be viewed largely as a struggle to resolve it (Dourish, 2001; McCullough, 1996).

Attempts to better match computational capability to human sensibility over the past twenty-five years tend to fit one of two perspectives, based on how they address the two worlds problem[1]. One approach has been to make digital spaces more hospitable for people by lending them familiar features of the physical world; the other has been to make computational capabilities more accessible from (and sensitive to) the physical contexts that people inhabit and share. Terminology for such approaches has not settled; throughout this thesis I will refer to the first approach as "Virtual" and the second as "Embodied". The term "Embodied" has various connotations in Artificial Intelligence, Philosophy and other disciplines; it is used here to describe an approach to designing computational systems that depends explicitly on physical forms and settings. While the Virtual/Embodied terminology is admittedly awkward, it is used here for historical reasons.

2.1.1. The Virtual Approach: Placing People in the Computer's World

One overall approach to improving human-computer interfaces is the "Virtual" approach: making digital spaces "inhabitable" by constructing an appearance of the familiar physical world. This approach is most clearly evident in "Virtual Reality" (VR) interfaces, but also characterizes interfaces that rely on a mouse and pointer pair.

2.1.1.1. Virtual Reality

Virtual Reality interfaces envelope the senses with computationally managed stimuli in a literal way in order to create the illusion of being *in* an alternate physical reality, a virtually-constructed world that is "real in effect but not in fact" (Heim, 1993)

---

[1] One more detailed taxonomy presented by (Dourish, 2001) distinguishes between four approaches; two will suffice for this thesis.

p109. While using a VR interface, stimuli from the actual physical world are replaced by artificial stimuli. A technical basis for exploring VR began with Ivan Sutherland's pioneering work in computer graphics and head-mounted displays at Harvard in the 1960s, however VR opened to widespread investigation and critique with Jaron Lanier's development of the "Data Glove", popularization of head mounted displays (as "Eyephones") and commercial offering of a system that combined these devices with software that allowed people to author their own virtual environments ("Virtual Programming Language") (Heim, 1993; Darken and Zyda, 2001). Presently the VR illusion is used towards numerous ends, from designing cars to discovering oil to developing new medicines (Heisse, 2003).

2.1.1.2. Mouse and Graphical User Interface

Literal envelopment of the senses is not strictly necessary for creating the illusion of inhabitable software. *Suggesting* a sense of the real has proven at least as powerful a means for allowing people to feel at home in software as *reproducing* a sense of the real. Since people sustain disbelief and "fill in the blanks", virtuality can (in the terminology of media theorist Marshall McLuhan) be served "hot" or "cool".

The mouse and mouse pointer are a particularly influential example of a "cool" virtual approach to addressing the two worlds problem of human computer interaction. Moving a mouse and watching its associated pointer doesn't look or feel *much* like having a finger in a virtual world, but it feels *enough* like this that people can sustain disbelief and work *as if* they do. The mouse was invented by Douglas Engelbart at the Stanford Research Institute in 1964, included with Xerox' "Alto", made commercially viable by Hovey-Kelly Design, popularized by Apple Computer with the introduction of the "Macintosh" in 1984, and is now available with most desktop computers (Pang, 2002). The graphical user interface, an idea that joined the mouse at Xerox PARC, further illustrates how mere allusion to physical things can facilitate comfortably acting as though *in* software; a GUI's representations of virtual rooms, desktops and trashcans make no pretense toward realism, however they are *enough* like their physical referents in appearance and behavior so as to suggest coherent courses of action to someone familiar with such things. By merely suggesting familiar aspects of physicality, a mouse and GUI effectively allow people to make themselves feel at home in software environments.

2.1.2. The Embodied Approach: Placing Computers in People's World

Despite the successes of the Virtual approach to facilitating human computer interaction, attempts to import salient aspects of the physical world into digital spaces have clearly left many things lost in translation, and many people wondering about the limitations of this approach. The "paperless office", once lauded as an ideal and a goal, is now widely dismissed as myth[2] and research in Human Computer Interaction (HCI) has turned increasingly toward the matter of "why matter matters" (Sellen and Harper, 2002; Weinberger, 2003).

In response to doubts about the Virtual approach and new reflections on traditional physical resources, an alternate perspective began to take shape. Pierre Wellner summarizes this shift in his article "Interacting with Paper on the Digital Desk":

> There is a difference between integrating the world into computers and integrating computers into the world. The difference lies in our perspective: Do we think of ourselves as working primarily in the computer but with access to physical world functionality, or do we think of ourselves as working primarily in the physical world but with access to computer functionality? Much of the research in human-computer interaction seems to emphasize the former perspective, yet many useful ideas can be gained from the latter. Instead of making us work in the computer's world, let us make *it* work in *our* world (Wellner, 1993) p.95.

Where the Virtual approach has emphasized *emulation* and often carried with it an implicit agenda of *replacement*, the emerging Embodied approach emphasized *augmentation* of physical things and spaces with computational capabilities, and tended to advocate the *preservation* of physicality. Where the Virtual approach is about evoking or simulating doorknobs, the Embodied approach is about enhancing their "doorknob-ness" (Negroponte, 1995).

Over the past twenty years, the Embodied approach has been manifest in a wide variety of commercial and research agendas, and has resulted in numerous products

---

[2] (in the sense of something believed but not true)

and intriguing demonstrations. The following sections describe a few of these, and discuss how they relate to each other.

2.1.2.1. Embedded Computation for Household Appliances

Before discussing more recent variants of the Embodied approach, it is useful to consider an older one: Embedded Computation for household appliances. While the desktop computer was making headlines, microprocessors were more quietly transforming the innards of everyday appliances for cooking, washing, mixing, cleaning and other household tasks. The forms and functionalities of these appliances were not radically changed by the microprocessor; chips were just cheaper, more reliable and more amenable to adjustment than the mechanism they replaced. The history of embedded computation for household appliances is less well publicized, however this particular Embodied approach to placing computation out in the world is worth highlighting for two reasons: First, interaction design has brought the paths of computer science and engineering together with the paths of industrial and product design, and second, the presently confusing climate of multimedia convergences and networked interactions is provoking renewed interest in appliances (Winograd, 1996; Sharpe, 2001).

2.1.2.2. Mobile Computing

Another venerable variant of the Embodied approach is Mobile Computing. The goal of Mobile Computing has traditionally been to make information and communication technologies smaller and more portable, so that they might be brought out into the physical settings where human activities traditionally take place. This paradigm’s success stories include the pocket calculator (first realized by Texas Instruments in 1967), the cellular phone (first prototyped by AT&T Bell Labs in 1977) and the Personal Digital Assistant (defined by Psion during the 1980’s) (Texas Instruments, 1995; Bellis, 2004; Long, 2000). Like Embedded Computation for household appliances, Mobile Computation has typically recreated previously existing functionality, albeit in much smaller packaging. Mobile Computation has, unlike Embedded Computation for household appliances, demanded new forms of physical expression.

2.1.2.3. Augmented Reality

The Embodied approach of Augmented Reality (AR) is about enhancing *physical* identities with *informational* capabilities (Mackay, 1996). In this approach, things and places are interpreted as base layers upon which additional informational layers might be overlaid. Or removed; at worst, according to this approach, physical things and environments must retain their valued physical characteristics; at best, they may additionally play host to new informational capabilities (Mackay, 2000). Augmented Reality is perhaps best described in terms of examples; two provocative demonstrations are the *DigitalDesk* by Pierre Wellner of EuroPARC in 1993, and *Aladdin* by Karon MacLean and Jayne Roderick of Interval Research in 1999.

The *DigitalDesk* was a physical desk with three unusual properties: 1) the capability of recognizing documents placed upon it, 2) the capability to project digital display onto these papers, and 3) the capability of responding to pointing gestures made over the desk's surface with a finger or pen. This set of capabilities allowed researchers at EuroPARC to explore accessing computational capability *through* physical paper (Wellner, 1993). In one explored scenario, sums could be calculated by simply by pointing at numbers in printed documents – no calculator required (Wellner, 1991).

While the *DigitalDesk* demonstrated *visual* overlay of a computational capability upon a physical resource, *Aladdin* illustrated the concept of informational overlay *haptically* (through touch). *Aladdin*, an ordinary doorknob equipped with some extraordinary interactive behaviors, was created to explore communicative potentials latent in threshold spaces (MacLean and Roderick, 1999). The programmed torque users experienced while turning *Aladdin* allowed several communicative scenarios to be explored, including "WellHouse", and "The Malleable Knob":

> <u>WellHouse:</u> As you step outside and close the door behind you, you wonder briefly if the lights in the living room were left on. The doorknob's smooth response as you leave reassures you. All is well, the house is at rest (MacLean and Roderick, 1999).

> Malleable Knob: "As people pass through the door and make use of the knob…levels of activity through the space are [reflected through it. The knob feels] firm and crisp at the beginning of [each] day [but grows] hot and gummy after bouts of heavy usage" (MacLean and Roderick, 1999) p6.

As *Aladdin* and the *Digital Desk* both illustrate, Augmented Reality is about situating informational capabilities within physical settings in a way that accentuates the value of those particular settings. As with Embedded Computing for household appliances, Augmented Reality tends to preserve physical forms and spaces. As with Mobile Computing, the value added by computation concerns information and communication. Augmented Reality has in common with "cool" Virtual approaches like desktop computing a metaphorical aspect, but this aspect is ultimately applied to things in the physical world rather than virtual entities.

2.1.2.4. Ubiquitous Computing

Ubiquitous Computing, Xerox PARC's agenda for computing research in the 1980's articulated by Mark Weiser and popularized through his article *The Computer for the 21st Century*, is based upon two observations: First, that "the most profound technologies are those that disappear" from the foreground of our attention and allow us to "focus beyond them on new goals", and second, that such "disappearing" technologies tend (somewhat paradoxically) to be everywhere (Weiser, 1991) p94. Weiser cites print and the electric motor as two examples. The printed word appears everywhere, but does not require our active attention. Similarly, there are more than twenty electric motors in a typical car, yet a driver rarely considers any of them[3]. Like the electric motor, the computer has evolved from big, expensive and centralized towards small, cheap and distributed, and the goal of Ubiquitous Computing has been to determine what the presence of computation everywhere should mean for the design of computational capabilities anywhere (Weiser, 1991).

---

[3] While much has been made of print, less has been made of the pencil as a "disappearing" technology. In "The Pencil: A History of Design and Circumstance" Henri Petroski notes with irony that while Henry David Thoreau seemed to think of everything when listing supplies for a pilot excursion in the Maine woods before Walden, he neglected to mention the pencil – even though he would have been lost in the woods without one, and made pencils by trade

While Ubiquitous Computing is much like Augmented Reality, there are a few differences. Ubiquitous Computing is not only about enhancing the identity of existing physical things and spaces informationally, it concerns new physical forms and spaces as well. Augmented Reality has typically demanded clear added value from computation up front, while Ubiquitous Computing has been more speculative in its introduction of computation "in the wild" (Gaver, 2002). While the motivations and assumptions of Ubiquitous Computing and Augmented Reality are somewhat different, the two agendas have, in practice, resulted in similar sorts of experimentation.

2.1.2.5. Tangible Computing

Tangible Computing emphasizes the expressive and subtle roles that touch and form can play in helping people to make sense of information and communicate a sense of human presence. Two examples characteristic of Tangible Computing include the *Marble Answering Machine*, a demo developed by Durrel Bishop at the Royal College of Art (RCA) in 1992, and *InTouch* a demo developed by Scott Brave and Andrew Dahley at the MIT Media Lab. *InTouch* consists of two identical sets of rollers, separated by distance but connected electronically. When the rollers of one set are rolled back and forth under the palm of one's hand, the rollers of both sets turn accordingly. Two people manipulating the sets of rollers in different places feel each others' presence, and can thus find ways to communicate through touch across distance (Brave and Dahley, 1997). The *Marble Answering Machine* illustrates the use of physical tokens (marbles) to represent digital information (telephone messages). By picking up a marble that rolls into a "messages in" tray and placing them in different hollows of a sculpted surface (the "telephone"), a user can hear the message repeated or return the call (Gamez et al., 2003). By abstracting the idea of communication through computation along tangible dimensions, *InTouch*, the *Marble Answering Machine* and similar demos have helped interface designers to glimpse some of the subtle, emotive and expressive potentials latent in the physical representation of information for communication.

Embedded Computation, Mobile Computation, Augmented Reality, Ubiquitous Computing and Tangible Computing all emphasize different advantages of physicality, however each of these variants of the Embodied approach shares the aim of situating computation in the human physical world rather than bringing people into computational spaces. While the older variants of the Embodied approach (Embedded and Mobile Computation) have clearly had an impact on everyday living and working, the gains of the more recent variants (Augmented Reality, Ubiquitous Computing and Tangible Computing) remain uncertain. Research based upon these technically challenging perspectives continues.

In the past section, two overall approaches to human-computer interaction have been described, and variants of each have been compared and contrasted. While the Embodied and Virtual perspectives are very different (in some sense opposites), they share the same over-arching goal: providing people with simultaneous and co-located access to the best of both worlds, physical and digital.

2.2. Physical-Digital Association

2.2.1. Motivation

Whether proceeding from an Embodied approach or a Virtual one, space matters immeasurably. From the simplest of switches to the most complex of immersive simulations, effective interfaces create and maintain reliable spatial correspondences between the sensed physical and the representational digital. As Weiser observes in *The Computer for the 21st Century*, “little is more basic to human perception than [spatial] juxtaposition” (Weiser, 1991) p95. Through spatial juxtapositions, we relate things to things, people to people, people to things, parts of one articulated body to another, and so on. Furthermore, spatial juxtaposition is deeply intertwined with categorization and identification; spatial relations make the difference between *this* and *that*.

While spatial relations and identities come as second nature to people, they are, by and large, inaccessible to software processes. This is true on a basic technical level, as well as on a subtler semantic level. On a technical level, microchips are ill equipped to sense physical distances, presences or differences. On a semantic level, computers lack access to human background regarding what these distances, presences and differences *mean*[4]. The present discussion addresses how the first of these “blindnesses” – the blindness computers exhibit towards physical identities and distances – may be reduced through application of sensing systems. First, this section introduces some of the technical concerns entailed in sensing physical identity and location; next it discusses one particular sensing technology in greater detail[5]. Identification and tracking are dealt with together in this discussion, as two aspects of the same technical challenge.

2.2.2. Tracking and Identification: Technical Considerations

There are numerous factors one must consider in choosing a sensor system for automated tracking and identification purposes, and in practice, these factors cannot be considered independently. No one sensor system works for all situations, and seemingly subtle differences in functional requirements can have a huge impact on what systems are appropriate. The list below presents some of the factors one must consider, and discusses several of their interdependencies.

<u>Origin</u> Location is a relation, and one must choose a frame of reference. The corners of a room, the earth’s magnetic field and an object’s position at an earlier moment are just some examples of references that might be used to anchor a coordinate system. Choosing any particular frame of reference has numerous implications for the technical solutions that are viable. For example, if the earth’s magnetic field is used as a reference, magnetic interference (caused by electrical currents and certain materials) may be an issue. If the corners of a room are used as references, “active” (powered) devices might need to be placed at these corners. If

---

[4] As McCullough notes in *Abstracting Craft*, “computers are very good as calculators, constructors… but [ultimately] we’re the ones who can see” (McCullough 1996). While the blindness computers exhibit towards human spatial and categorical sensibilities is extremely important, and a major source of confusion in human computer interaction, it is not the subject of this section and is beyond the scope of this thesis. For a good treatment of this topic, see *Understanding Computers and Cognition* by Winograd and Flores (1986).

[5] More comprehensive discussion relating specific technologies to each other can be found elsewhere (Paradiso, 2000; Hightower and Borriello, 2001; Leydon, 2003).

an object's former position is used as a reference, measurement errors may accumulate over time.

Degrees of Freedom Tracking can be carried out to varying "degrees of freedom". In some situations, proximity – radially how "close" or "far" something is – may be sufficient. In other situations, such as tracking the position of a computer mouse, (Cartesian) two-dimensional relative position is sufficient. More demanding situations like sculpting virtual clay, may benefit from (Cartesian) three-dimensional tracking of position as well as orientation.

Marked vs. Unmarked Sometimes the important differences between things are detectable and physical (e.g. metal items passing through a metal detector). In other situations, key differences may be neither physical nor detectable (e.g. ownership of identical copies of an item by different people). These two types of situations lend themselves to two different strategies for automatic identification (autoID): unlabeled and labeled. Unlabeled identification relies on the sensed properties of things themselves for discriminatory power, while labeled identification relies on the sensed properties of markers *attached* to the things themselves. Labeled identification entails the overhead of labeling, but simplifies identification (since labels are specifically designed to be identified). Labeling makes setting an item's "visibility" to software an explicit choice, but introduces an added potential for confusion: unplanned separation of label from item.

Sensed Phenomenon A variety of phenomena can be automatically sensed. Color, pressure, capacitance, (galvanic) electrical contact, weight, distance, temperature and magnetic field strength are just a few possibilities. Choosing to sense one phenomenon impacts the sorts of objects that might be recognized, and places limitations on the sorts of environments where recognition may take place. Choosing to sense color, for example, impacts the sensing of objects that are visually occluded or in low light environments.

I.D. Resolution The discriminatory power of an identification system dictates how many different (sorts of) items can be automatically detected. The Electronic Article

Surveillance (EAS) systems commonly used to prevent shoplifting of CDs and other retail items have one-bit resolution; tagged items leaving a shop are recognized as either legal or illegal[6] (Hsiao, 1999). The (UPC) bar codes used in supermarkets are capable of distinguishing between billions of kinds of products (Lefferts, 2002).

Spatial Resolution The accuracy and precision with which spatial relations are reported vary greatly from one tracking system to another. Some systems present spatial information in a purely discrete fashion (e.g. "passing through door #1") while others provide effectively continuous measures. One continuous system may report measurements on the order of millimeters, another on the order of meters.

Range A concern closely related to spatial resolution is range: the maximum distance at which things can be sensed. Some sensing systems (such as those in touch screens and certain light switches) may have a maximum range on the order of a millimeter or less; the global positioning system (GPS) tracks effectively around an entire planet.

Contact Requirements Physical contact is prerequisite for some sensor systems. Pressure sensitive floor tiles, for instance can only be used to locate and identify what weighs upon them. Other sensor systems, such as those for reading "smart" card IDs, require electronic as well as physical contact for identification. Still other systems, such as the Radio Frequency Identification (RFID) tags typically used for inventory management in warehouses, require neither physical nor galvanic electrical contact.

Temporal Resolution Some sensor systems (such as a mouse) operate quickly enough to provide feedback in "real time", i.e. without perceptible delay. Other sensing systems (such as high resolution computer-vision based motion capture systems) entail noticeable delay. In some cases, temporal resolution may be "traded off" to improve spatial resolution (as is the case with GPS) or I.D. resolution (as is possible with RFID).

---

[6] Strictly speaking, either an EAS label is recognized or *no* EAS label is recognized.

Number of Objects In some systems, the presence of one item within range precludes identification of other in-range items. Alternatively, a sensor system may track and identify multiple items simultaneously. Anti-shoplifting systems don't discriminate between one or more shoplifted items that pass through a doorway simultaneously; marathon timing systems, on the other hand, must identify multiple finishers as they pass over the finish line in close succession.

Line-of-Sight Requirements Some sensor systems (such as bar code systems) require an unobscured path, a "line-of-sight", between sensors and sensed in order for identification to take place. Other systems (such as RFID) have no such line-of-sight requirement.

Power Requirements Powered sensing systems may allow greater range, spatial resolution, temporal resolution or discriminatory resolution than systems that don't require power, but these advantages come with strings attached – literally, in the case of power cables, and figuratively, in the case of batteries. (Batteries die, require recharging and are composed of particularly toxic heavy metals).

Privacy With the territory of identification and tracking come numerous privacy issues. Who should have license to automatically identify and track, in what situations and toward what ends? Should the "visibility" of tracked physical entities to software be changeable from these entities, from software, or from both? Who should have permission to change this visibility? If identification requires broadcasting data, can third parties listen in upon this transmission? If so, under what conditions is this acceptable?

Cost The cost of any required sensors and labels impacts the viability of a tracking solution. For example, the current cost of RFID labels (on the order of 10 cents each) places them out of reach as a viable way to associate physical items with digitally managed prices in supermarkets, but makes them a cost-effective means to track palettes and crates in the shipping industry.

While the above considerations neither form a comprehensive list nor provide a systematic account of specific sensing methods' strengths and weaknesses, they serve to give a flavor for the technical challenge inherent in maintaining physical/informational correspondences. This challenge must be acknowledged and addressed at some level by most human computer interfaces.

2.2.3. Magnetically-Coupled Passive Resonant Digital ID Tags

2.2.3.1. Overview

While no one sensing modality can facilitate tracking and identification under all conditions, Radio Frequency Identification (RFID) has emerged as one of the more versatile and robust solutions available. This section describes RFID tracking in terms of the considerations outlined in the previous section, explains RFID's underlying mode of operation, and introduces a latent potential of RFID that might be tapped; an unused aspect of particular relevance to the work described in this thesis.

Radio Frequency Identification is a labeled approach to facilitating physical/digital correspondences. An RFID system typically consists of one fixed "interrogator" (also called a "reader/writer") and numerous mobile "transponders" (also called "tags" or "labels") that can be embedded in or affixed to the surface of physical items to be identified. Positioned at checkpoints, interrogators report on the presence or absence of transponders with discrete spatial resolution (i.e. "Transponder X is present in doorway Y"). For an illustration of an RFID system, see Figure 1.

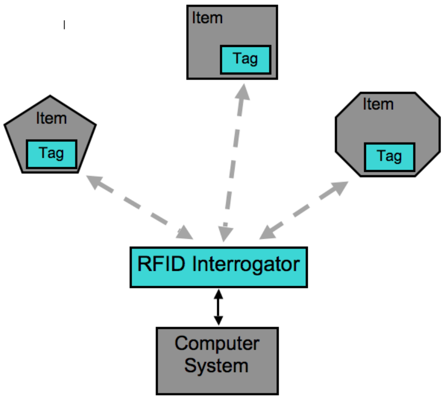

Figure 1: High-level block diagram of an RFID system; an interrogator relays a computer's "read" and "write" commands to tag transponders affixed to physical items. Tag ID numbers are sent back, through the interrogator, to the computer system.

An RFID interrogator remotely powers and communicates with one (or more) in-range transponders, and relays transponder data to a “host” computer system where it becomes accessible to software. Transponder data typically consists of an ID number and several programmable fields, all digitally represented. Transponders are powered, identified, read from and written to without physical or galvanic electrical contact, through the phenomenon of electromagnetic resonance. Some (low frequency) RFID systems rely on “near field” inductive coupling in the manner of an electrical power transformer, while other (high frequency) systems rely on “far field” reflective coupling, like a RADAR system (Finkenzeller, 2002). Since tags are powered by the same electromagnetic field that enables communication, they require no power cables, batteries or power “scavenging” subsystems. Because tag and reader are electromagnetically coupled, unobstructed line-of-sight is not required for identification. Electromagnetic fields can be influenced by the presence of metal, thus metal objects in the vicinity of an RFID system may have an effect on the system’s read range.

ID resolution, temporal resolution and range all vary from system to system, however systems capable of distinguishing between 2^{64} possible ID numbers more than 25 times a second at ranges on the order of a meter are increasingly common (Texas Instruments, 2001; Texas Instruments, 2002). The price of interrogators for such systems ranges from a few hundred to a few thousand dollars, while the cost of transponders for such systems ranges from twenty-two cents (bulk pricing) to one dollar (small sample pricing) (Wallington, 2004).

One indication of RFID’s technical success as a bridge between physical entities and digital representations is the fact that its application is now limited as much by privacy concerns as by feasibility. News articles on RFID run under such headlines as “Big Brother in Small Packages”, and several companies have been forced to abandon plans to tag merchandise in the face of boycotts by consumer-advocacy groups and widespread suspicion (McCullagh, 2003; BBC, 2003; Dougherty, 2003).

The full extent to which RFID will serve as an *acceptable* means of coordinating the physical with the digital remains to be seen.

2.2.3.2. RFID: Principles of Operation

While RFID facilitates automated identification and location in many situations, it fails to do so in other situations and its failure modes can be frustratingly inscrutable. In order to understand the limitations and tradeoffs that accompany RFID tracking, it is first necessary to understand how RFID functions. This section introduces the operating principles of the most common type of RFID system: the inductively coupled, low frequency system (Finkenzeller, 2002).

At the basis of RFID is the resonator: a device that oscillates in response to an excitation. A resonator vibrates most strongly in response to an excitation at its "resonant" or "characteristic" frequency f_0. As the frequency of excitation moves above or below f_0, the resonator's intensity of vibration decreases accordingly. The steepness of this decrease is described by the resonator's "quality factor" Q. "High-Q" resonators respond selectively to a narrow range of frequencies around f_0, while "Low-Q" resonators vibrate in response to a wider frequency range surrounding f_0. A higher quality factor is desirable for effective power transfer from exciter to resonator, while a lower quality factor ensures predictable resonant response in the face of real-world tolerances. Figure 2 shows the frequency response of two resonators with the same resonant frequency, but differing quality factors.

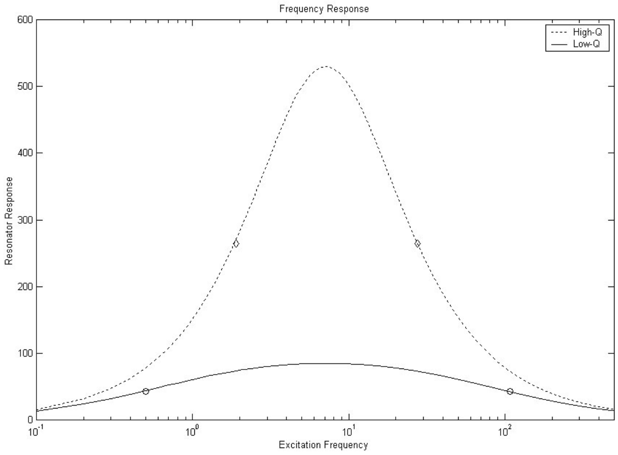

Figure 2: This graph shows the frequency response for two resonators with the same resonant frequency but different quality factors. The high-Q resonator's amplitude of vibration *at* the resonant frequency is highest, while the low-Q resonator's amplitude of vibration *around* the resonant frequency is less variable. (The two selected points on each curve indicate 50% drops in signal amplitude from their maximum values. Notice the horizontal distance between 50% points, the "bandwidth", is wider for the low-Q curve).

Electronically, resonators can be modeled using three discrete circuit elements. A capacitor (a device that passes high frequency electrical currents while blocking low frequency currents) combines with an inductor (a device that does the opposite) to set a resonator's resonant frequency f_0. With capacitance and inductance values set, the value of a resistor (a device that resists the flow of electrical current in a frequency-independent fashion) determines Q. For a circuit diagram of a parallel resonator, see Figure 3.

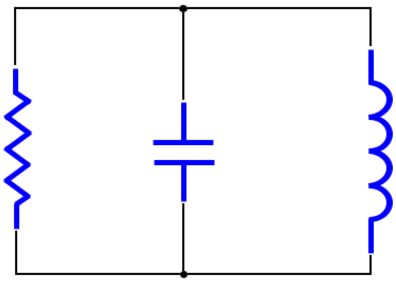

Figure 3: A parallel resonator, composed of a resistor (left) a capacitor (middle) and an inductor (right).

Inductive and capacitive circuit elements are practically realized within an RFID tag as a spiral of thin metal foil (inductor) and a pair of foil patches separated by a plastic substrate (capacitor). Sources of resistance in a real RFID tags include parasitic resistance in the foil material, and load resistance due to the tag's transponder chip. For a picture illustrating the practical realization of a resonator as an RFID tag, see Figure 4.

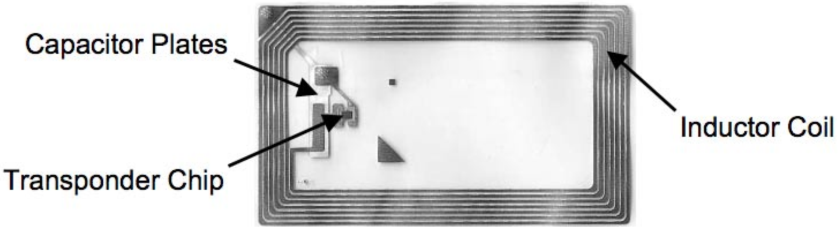

Figure 4: An RFID Tag. On the perimeter is the foil coil inductor, to the left are the foil plate capacitor and transponder chip.

When a switch is placed in the circuit, as shown in Figure 5, the resonator can be "tuned" and "detuned"; closing and opening the switch starts and stops resonation. The difference between a tuned resonator and a detuned resonator can be sensed remotely at the interrogator's antenna as a slight change in signal amplitude. This remotely detectable difference between tuned and detuned resonators is sufficient for enabling the Electronic Article Surveillance (EAS) systems that are used to prevent the shoplifting; when EAS tags are deactivated at check-out, the cashier is effectively detuning a tuned resonator.

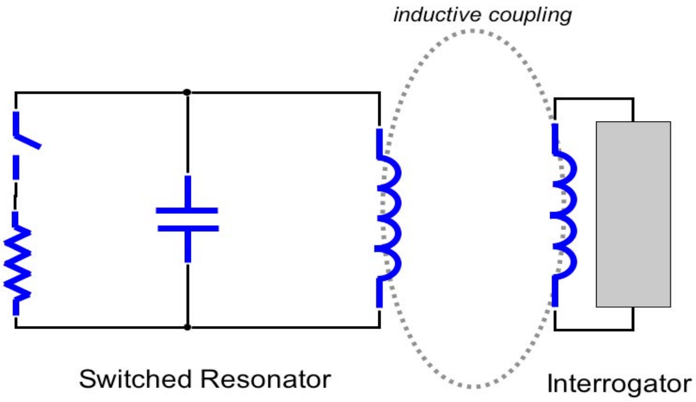

Figure 5: A switched parallel resonator inductively coupled with a remote interrogator. When the switch is open, resonator is tuned and the interrogator can sense its presence. When the switch is closed, the resonator is effectively detuned and made invisible to the interrogator.

Radio Frequency Identification takes the remote sensing of tuned and detuned resonators one step further; by alternately detuning and tuning the resonator at precise intervals through means of an electronic switch (transistor), RFID tags remotely modulate the signal amplitude of the interrogator's antenna. This modulation can be detected and decoded by the interrogator as the binary 1s and 0s that comprise the tag's ID number. For a diagram illustrating this process, see Figure 6.

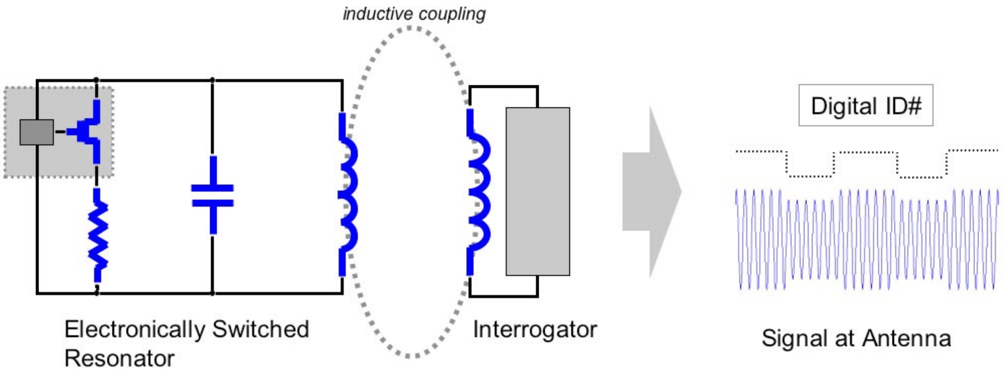

Figure 6: In an RFID system, a parallel resonator is actively tuned and detuned by a transistor switch (the transistor is driven by the tag's transponder chip). This repeated tuning and detuning remotely modulates a signal at the interrogator's antenna (the signal that powers tag resonation). The modulated signal is then decoded to reveal the tag's digital ID number.

In summary, the process is as follows: Oscillating current in the interrogator's antenna causes the tag's resonator to resonate through inductive coupling. Resonation charges a reserve for switching the tag's transistor on and off in a predefined pattern unique to the tag. The transistor is switched on and off repeatedly to alternately tune and detune the tag. The interrogator remotely senses this tuning and detuning as slight changes in signal amplitude at its antenna, and this modulation is detected and decoded[7] to reveal the tag's digital ID.

In light of the above explanation, it is now possible to discuss some of the many tradeoffs inherent in RFID system design.

<u>Range vs. Tag Size:</u> Increasing the size of a tag can increase the range of an RFID system; larger tags can embody larger inductors and can thus capture a greater portion of an interrogator's magnetic field. As tag size increases, however, tags become unwieldy as labels for smaller physical items.

<u>Tag Distance vs. Tag Access Frequency</u>: As a tag is moved away from the interrogator and out towards the maximum read range, it couples less and less strongly with the interrogator's magnetic field. As a result, the time necessary to charge the reserve that powers switching the tag transistor on and off increases, and the number of times the tag can report its ID each second decreases.

<u>Range vs. Robustness</u>: Tags that embody High-*Q* resonators couple efficiently with an interrogator and thus enable greater range than tags embodying low-*Q* resonators. Unfortunately, high-*Q* tags are often impractical; they can be more expensive and difficult to manufacture, and are easily detuned (that is, shifted to a slightly different resonant frequency) by nearby metal items – including other tags. Tags that embody low-*Q* resonators are less efficient, but more robust, cheaper, and easier to manufacture.

---

[7] There are a number of techniques for encoding/decoding a tag's digital signal; they are not ultimately relevant to this thesis, and so they will not be discussed.

Number of Readable In-Range Tags vs. Tag Access Frequency: Since all tag IDs share a common resonant frequency, "collisions" between tag transmissions can be a problem. Anti-collision algorithms can effectively eliminate this problem (by forcing tags to take a number and wait in line) but introduce delay. Since tags all share a common frequency and transmissions are handled serially, the frequency with which a given tag can be identified decreases as the number of in-range tags increases.

Range vs. Safety, Legality and Implementation Complexity: While RFID systems operating at higher resonant frequencies offer extended range, high frequency (i.e. microwave) operation is accompanied by more stringent legislation and health concerns. Additionally, designing and prototyping at high frequencies is significantly more complicated; the independences and simplifications of low-frequency electronics begin to break down and measurement becomes more challenging.

ID Resolution vs. Transmission Time: Because IDs are serially transmitted, there is a tradeoff between ID resolution and transmission time. Since transmission time is short (typically on the order of a few milliseconds) and ID resolution grows exponentially with the number of bits used to encode an ID number, this tradeoff is typically not a limiting factor.

These are just a sample of the many tradeoffs inherent in RFID system design; they are included here to communicate a sense of the multiple, simultaneous high-wire balancing acts that have been successful each time an RFID system makes an item's physical presence automatically accessible to software.

2.2.3.3. Untapped Potential for Continuous Spatial Resolution

While RFID systems demonstrate exquisite ID resolution (2^{64} ID numbers could easily account for all the grains of rice eaten on earth over the course of a year) their spatial resolution is somewhat less spectacular (Ericson, 2002). Typically, tag proximity is reported with mere binary resolution; a tag is either present within the vicinity of an interrogator, or it is absent[8].

The binary spatial resolution uniformly provided by commercially available RFID systems masks a potential for higher spatial resolution that is latent within RFID communication. Since the strength of an electromagnetic field varies with distance, the received signal strength of a tag transmission might be used to estimate the distance between tag and interrogator. The graph in Figure 7 depicts a measured relationship between (axial) distance and magnetic field strength reported by Galbourne (2003).

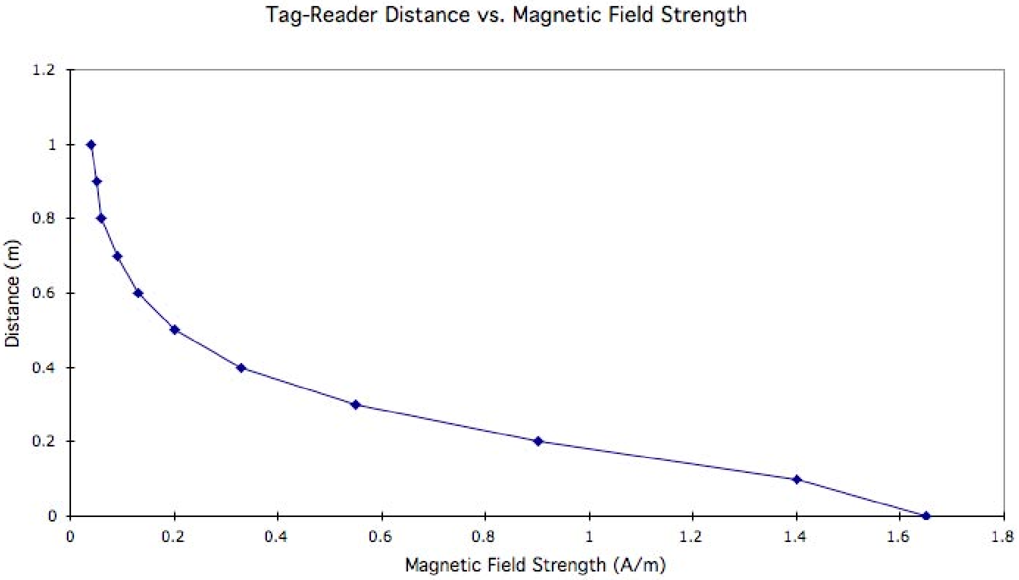

Figure 7: Estimated interrogator distance as a function of magnetic field strength.

Since distance can be treated as a function of magnetic field strength, it is *in principle* possible for an RFID interrogator to report an estimate of tag proximity along a continuous scale.

Exactly *how* tag transmission strength varies with distance is not a trivial matter. It depends, for one thing, on the size and geometry of the interrogator's antenna and the

---

[8] The fact that operational amplifiers are expensive and slow, while digital comparators are cheap and fast helps to explain why commercially available RFID interrogators report so universally with binary spatial resolution.

tag's inductor. For another, different mathematical descriptions apply for when a tag is close to an antenna than for when a tag is far from an antenna. In addition, nearby tags are likely to influence the strength of each other's signal transmissions. Furthermore, magnetic coupling between interrogator and tag is not only a function of relative distance, but also a function of relative orientation. Figure 8 illustrates this dependence on orientation; while tags A and B are approximately the same distance from the interrogator's antenna, tag A's perpendicular orientation facilitates magnetic coupling, while tag B's parallel orientation suppresses it.) To exacerbate these complications, the fact that magnetic field lines divergence rapidly with distance means that the precision of distance estimates based upon magnetic field strength will diminish with distance as well. Deriving distance estimates from magnetic field strength is clearly a difficult technical challenge, however improving upon the current benchmark of binary spatial resolution is clearly possible.

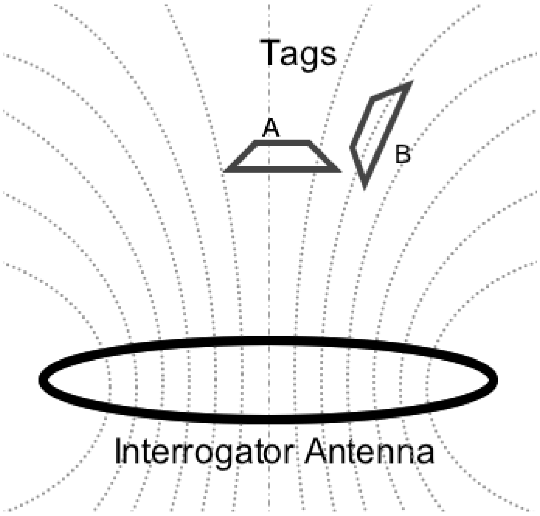

Figure 8: The relative orientation of interrogator antenna and tag inductor impacts the strength of inductive coupling. Tag A is parallel to the interrogator antenna, and thus perpendicular to the interrogator's magnetic field lines. As a result, this tag is strongly coupled with the interrogator. Tag A is perpendicular to the interrogator, and thus parallel to the interrogator's magnetic field lines. This configuration results in weak or non-existent coupling.

Although continuously tracking the proximity and location of passive resonant tags based on magnetic field strength is a difficult analytical and technical challenge, it has already been demonstrated to function. For his master's thesis at MIT, Kai-Yuh Hsiao developed several tag-reading devices that mapped received signal strength to spatial information (Hsiao, 1999). One of these devices, based on a single antenna coil, reported a rough estimate of tag proximity sufficient to drive a musical controller demonstration (see Figure 9, left side). Another more ambitious device, based on three mutually perpendicular pairs of antenna coils, reported the position and orientation of a cluster of tags in three-dimensional space. This prototype resolved position in three dimensions to within +/- 2cm, and orientation along three axes with a precision of approximately +/-10° within a cubic volume of approximately 30cm dimensions (see Figure 9, right side). A larger version of this second device is currently under construction.

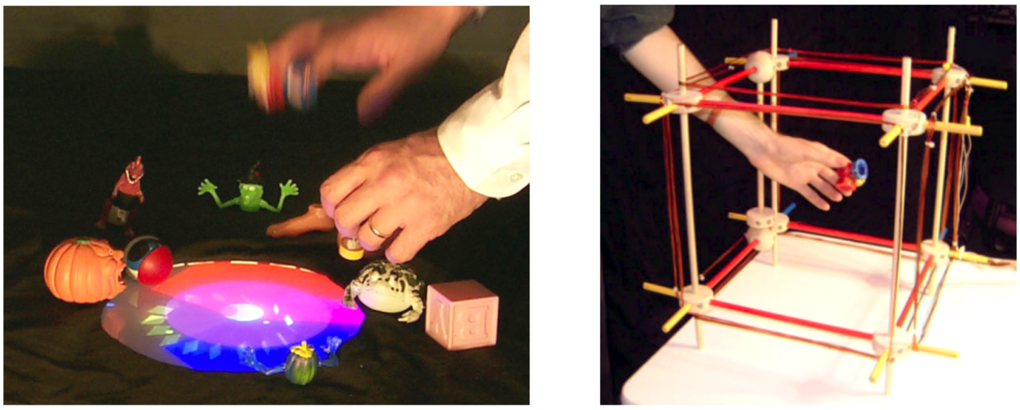

Figure 9: "SweptRF" tag tracking systems. Left: One coil system for tracking tag proximity. Tags (embedded in small plastic toys) can be moved above a table surface (containing a tag reader) to control a musical application. Right: Three orthogonal coil-pair system for tracking position and orientation. A cluster of tags is embedded in the hand-held device; coil-pairs circumscribe opposing cube faces.

The tag tracking devices created by Hsiao demonstrate the feasibility of estimating tag distance based on received tag signal strength. The resonant tags used in these projects were EAS tags with low ID resolution (tags were identifiable by frequency) rather than digital RFID tags with high ID resolution, however the overall approach is applicable to RFID tags as well. While typical RFID systems only support binary spatial resolution, they hold an untapped potential for continuous spatial tracking.

Maintaining correspondences between physical entities and digital representations is a challenge interface designers repeatedly confront, and sensor systems provide the means for making physical identities and spatial relations accessible to software. The previous two sections have introduced several technical considerations related to tracking and identification then discussed one particular technology, RFID, in the light of these considerations. The tradeoffs inherent in RFID, its principles of operation and its untapped potential to provide continuous spatial resolution have all been discussed.

2.3. "Focal Gradients" and Orienting Toward Tools for Mediated Display

In efforts to develop better computational tools, it can be instructive to first reflect on the strengths and weaknesses of traditional tools, tools with a longer (and perhaps richer) evolutionary heritage than the personal computer. Here we consider several traditional tools for extending human perception, identify a valuable feature that they share, and examine how this feature has so far been applied within the context of digital tools.

Many of the tools we use serve to extend our perceptions so that we can see, hear and feel things that are not perceptible through our senses alone[9] (McLuhan, 1964; McCullough, 1996). Flashlights extend our sight into the dark. Magnifying glasses allow us to see on a much smaller scale. Through a radio, we hear music from other places and other times. A Geiger counter allows us to hear the presence of radiation, mercifully before its effects can be directly sensed. Flashlight, magnifying glass, radio and Geiger counter are a few of the numerous timeworn tools we rely upon for mediated (indirect) perception.

While traditional tools for mediating perception are diverse – extending different senses towards different ends through different physical means – they have this in common: traditional tools we employ to reveal hidden features of the physical world provide feedback with a focus. When we move a magnifying glass horizontally over a subject of examination, the subject appears largest when under the magnifier's center. When we aim a flashlight at a target, the target appears brightest when near the center of the flashlight's beam. Tuning a radio through its knob amounts to finding strong sonic centers amidst weaker sonic peripheries. When we wave a

Geiger counter closer and closer to a source of radiation, it clicks with gradually increasing frequency. Traditional tools for mediating perception tend to provide a *focal feedback gradient*: a stimulus that declines continuously from a strong centre towards a weaker periphery[10].

The focal gradients provided by traditional tools for mediating perception serve a useful function: they guide us into alignment with phenomena we seek to examine. Through such gradients, we orient ourselves with respect to tools and subjects of observation. In engineering terms, this type of dynamic alignment is effectively modeled as a system approaching steady state behavior through negative feedback, with such characteristics of travel time, overshoot and settling (see Figure 10).

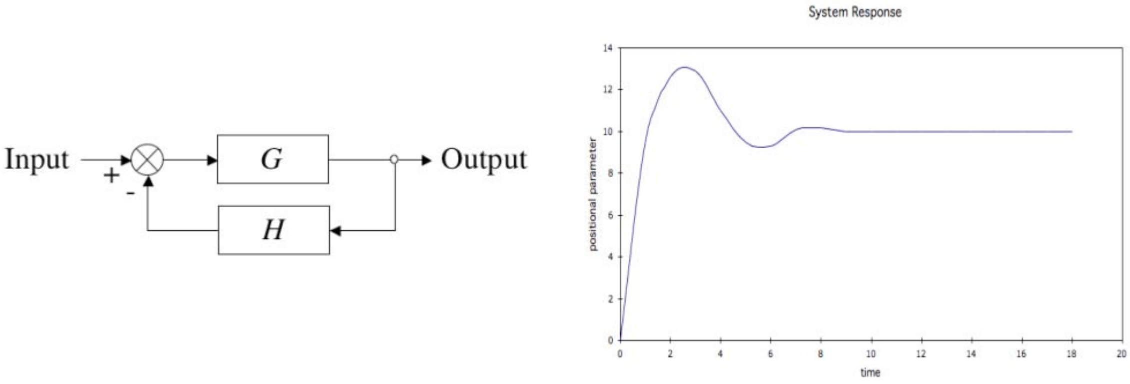

Figure 10: Left: A generic block diagram of a dynamic system with negative feedback. Right: A generic graph showing a parameter traveling towards, overshooting and finally settling at a steady state through negative feedback.

In addition to promoting effective alignments, focal feedback gradients serve another purpose. On a much subtler level, they contribute to our sense of seamlessness as we move between direct and mediated perception. Consider the apparent size of the magnified mouse in the sequence of frames in Figure 11. As the mouse moves closer and closer to the periphery of the magnified region, its apparent size grows closer and closer to its size *without* magnification so that there is no "jump" in apparent size

---

[9] McLuhan's discussion of "media" is sufficiently broad so as to encompass "tools".

[10] Whether this is due to design, evolution or the nature of physical phenomena is open for debate. The present discussion can continue regardless.

when the mouse moves out from under the glass lens. The focal gradient of a flashlight beam facilitates a related seamlessness; as the target begins to exit the flashlight's beam, it begins to appear as it will ultimately appear: unlit by the flashlight. Similarly, moving a Geiger counter away from a source of radiation leaves its wielder in a position perceptually similar to not having the Geiger counter's assistance. If static is disregarded (and if stations are not tightly packed) "tuning out" with a radio knob gradually brings a listener into the situation of not listening through the radio. The focal gradients in feedback that traditional tools provide help us to transition seamlessly between spaces of direct and mediated perception.

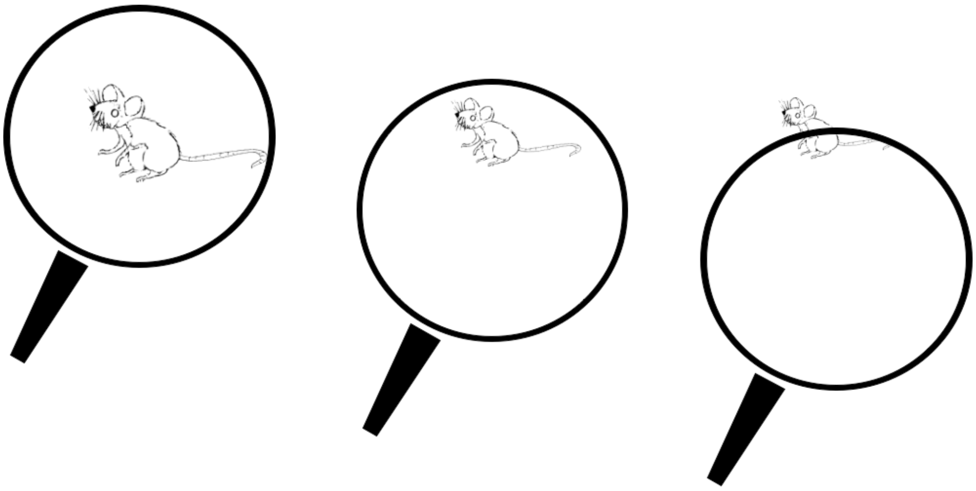

Figure 11: As the mouse moves away from the centre of the magnifying glasses field, its apparent size approaches its actual size. This decline in magnification enables a seamless visual transition at the glasses edge, between mediated perception *through* the glass, and direct perception *without* the glasses aid.

In transitions between the physical world and *digitally* mediated spaces, on the other hand, seamlessness represents a fundamental challenge. As Hiroshi Ishii and Brygg Ullmer lament in *Tangible Bits*, "The absence of seamless couplings…leaves a great divide between the worlds of bits and atoms" (Ishii and Ullmer, 1997) p234. The "seamfulness" of digitally mediated display – in contrast to the seamlessness of more traditional tools – leaves users without clear means for continuously maintaining and adjusting alignment between self, tool and subject of examination.

In numerous settings, the focal gradients provided by traditional tools are now being emulated to bring seamlessness to people's experience of digital tools. Research in force feedback has suggested the value of "attractive basins" "magnetic buttons" and "gravity wells" for guiding hands toward useful features of virtual spaces (Langdon et al., 2002) (Dennerlin et al., 2000). "Hyperbolic" visual displays that synthesize magnification around a mouse pointer effectively increase screen "real estate"[11] (Lamping and Rao, 1996). Focal gradients in audio amplitude facilitate finding music within "sonic browsers" (Brazil and Fernström, 2003). Although digital tools for mediating perception may be based upon abruptly changing binary representations, the feedback they provide need not "feel" digital. Through synthesizing focal gradients characteristic of more traditional tools, it is possible to bring a degree of seamlessness to transitions in digitally-mediated display.

2.4. Haptic Display

While computer interfaces have grown steadily more accommodating of our visual and auditory senses over the past 40 years, they continue, by and large, to make poor use of our rich sense of touch and own-body alignment, and thus attenuate a powerful (if quiet) way of knowing. *Haptics* (touch, together with kinesthetic sense) closely couple sensing with manipulation, and many regard it as a "missing piece" in human-computer interaction. This section discusses some of the motivations and means for computationally controlled active haptic display.

2.4.1. Motivations for Haptic Human Computer Interaction

Presently, computer interfaces respond poorly to our haptic sense. Our bodies hold the expressive potential of 360 joints; our skin can sense touch, deep pressure, vibration and temperature, however these extensive kinesthetic and tactile dimensions of perceptual experience go largely untapped in our interactions with computer systems (Lloyd and Sivin, 2002; Klatzky and Lederman, 2002). Indicative of interfaces' haptic insensitivity is the scant attention that many interface designers

---

[11] The "Dock" within the Mac OS X Finder illustrates the principle of hyperbolic visual display.

pay to the skill and experience of human hands. As Malcom McCullough observes in *Abstracting Craft*:

> Hands are underrated. Eyes are in charge, mind gets all the study, and heads do all the talking. Hands type letters, push mice around, and grip steering wheels, so they are not idle, just underemployed. This is a sorry state of affairs, for hands contribute much to working and knowing. By pointing, by pushing and pulling, by picking up tools, hands act as conduits through which we extend our will to the world. They serve also as conduits in the other direction: hands bring us knowledge of the world. Hands feel. They probe. They practice. They give us sense, as in good common sense, which otherwise seems to be missing [from interfaces] lately (McCullough, 1996).

Clearly, the knowledge embodied in human hands is extensive, if under-appreciated.

While hands are not the *only* instruments of our haptic sense, they are its primary focus, and marginalization of the hand in particular has gone hand in hand with failing to accommodate the body's rich haptic sense in general. Reflecting on the disconnection between haptic and visual feedback promoted by the graphical user interface, PARC researcher Ranjit Makkuni draws far reaching conclusions: "we missed a fundamental connection: the integration of the hand and the eye in the act of interaction. The GUI form ultimately disembodied the learner, in turn, creating the static office" (Sullivan, 2002) p59. This "disembodiment" – a numbness towards the haptic sense – is what Douglas Coupland's archetypical programmer expresses in *Microserfs*: "I feel like my body is a station wagon in which I drive my brain around, like a suburban mother taking the kids to hockey practice" (Coupland, 1996) p4. Computer interfaces' poor accommodation of human haptic sensibilities have a powerful impact on how we *feel* while using computer systems.

The desire to restore tactile and kinesthetic dimensions that were lost in the initial rush to embrace digital tools is one motivation for applied haptics research. Another is the increasing number of small Mobile and Ubiquitous devices people use while on the go and in public spaces. Pagers, mobile phones, digital cameras and personal digital assistants all lack the surface area for extensive visual display, and their mobile use makes loading visual attention undesirable. The extent to which sound can serve as an effective communication channel is further limited by the fact that such devices frequently occupy and share public settings, where privacy and social interruption are concerns. Simultaneous requirements for small size, low visual loading, minimal interruption and privacy make haptic feedback attractive as an alternative or additional channel for communication.

2.4.2. Means for Haptic Communication

Since the haptic sense is a composite sense encompassing proprioception (kinesthetic own-body sense), light touch, deep pressure, vibration, heat and pain, there are in principle many perceptual dimensions along which stimuli might be modulated, individually or together, to communicate information. Practically speaking, two approaches have proven particularly useful: active force feedback (AFF), and vibrotactile feedback.

Active force feedback synthesizes forces and torques and delivers them to the body so that one can push prod and poke at virtual things – or manipulate "real" physical things – at levels of strength, distance, or sensitivity not possible through (literal) direct manipulation alone. Availing of active force feedback typically requires holding onto or being strapped into an articulated electromechanical linkage. As one moves, the linkage tracks one's movements and delivers spatially appropriate forces and/or torques from virtual or remote subjects of manipulation. (For a picture of one apparatus for active force feedback, see Figure 12). Active force feedback allows people to act where they cannot easily go (e.g. nuclear reactor cores, outer space, the deep sea, the human body), invoke machine powers with human sensitivity (e.g. exoskeletons explored by the military), and apply craft skills within a notational context (e.g. sculpting virtual clay to create "organic" models that can be printed in 3D, saved, and copied) (Stone, 2000; Novint, 1999; McCullough, 1996). Active force feedback closely couples sensing with manipulation and holds the potential to support "precise and rapid controls" (Verplank et al., 2002) p33. Because active force feedback entails tracking motions and delivering large forces with high spatial and temporal resolution, AFF systems tend to be large, heavy, costly, computationally intensive, power-hungry and mechatronically complex in comparison with other sorts of input/output devices. These characteristics have so far limited AFF to costly stationary systems used within highly specialized application domains.

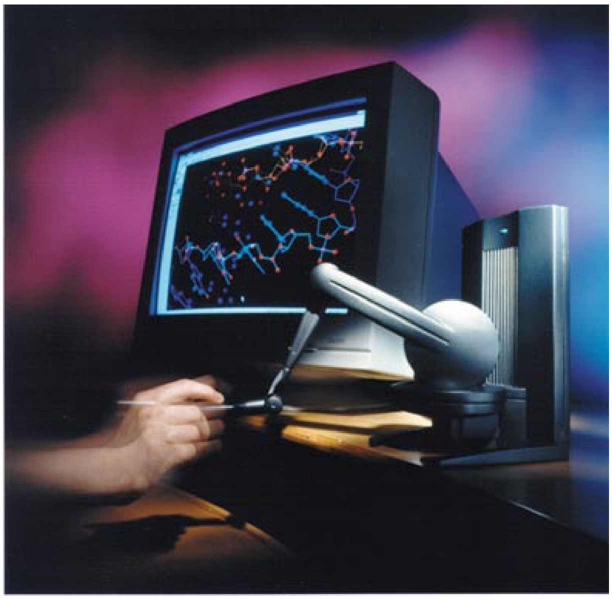

Figure 12: The PHANTOM® Desktop™, an active force feedback device that allows users to experience artificial forces while moving a pen-tool in three-dimensional virtual space. (Photo courtesy of SensAble Technologies, Inc.).

Systems incorporating vibrotactile feedback, in contrast to active force feedback systems, tend to be small, light, cheap, efficient, computationally un-intensive and mechatronically simple to realize. These technical and economic advantages – together with the previously mentioned constraints intrinsic to Ubiquitous and Mobile devices – have prompted widespread interest in vibrotactile haptic display.

Numerous recent studies on vibrotactile feedback build upon earlier work on the psychophysics of touch (Verrillo and Gescheider, 1992) (Kenshalo, 1978). One recent study attempts to "distill underlying perceptual dimensions of vibrotactile stimuli" (Erp, 2003) p111. Two others concern synthesizing one perceptual dimension, "intensity", from two physical ones, amplitude and frequency, to increase the effective dynamic range of a vibratory stimulus (Murray et al., 1998; Murray et al., 2003). Psychophysical results have been collated and packaged to inform design – Erp (2002) provides "guidelines for tactile display", while Challis and Edwards (2000) offer "design principles for tactile interaction" – and a new breed of haptic designers are already pushing the envelope. Chang et al. (2002) have explored vibrotactile feedback's potential for mediating interpersonal communication, while

Gunther et al. (2002) have put forward felt vibration as a rich (cutaneously groovy![12]) compositional medium in its own right. On a technical level, new tactile actuators are in the pipeline; Pasquero and Hayward (2003) are developing the vibrotactile equivalent of pixels that enable the presentation of vibration differentially across a surface with 1mm resolution. On perceptual, technical and application levels, vibrotactile display is a growing research interest.

Active haptic feedback offers numerous advantages for interaction. In the context of stationary systems, active force feedback extends manual reach into previously inaccessible environments and accommodates manual skill for working with computers to an unprecedented degree. Vibrotactile feedback alleviates visual loading, supports private communication and reduces public interruption in mobile contexts of use.

2.5. Fitts' Law for Targeting

A perennial source of consternation for those who would see human-computer interaction design as hard science is the difficulty of descriptive and predictive mathematical models to formulate a convincing footing. The sheer number of factors influencing people's interactions with and through computation, the reality that many salient details are practically unmeasurable, and the inherent tradeoff between tight experimental control and natural, situated interaction have all conspired to limit the effectiveness of mathematical models for human-computer interaction.

One mathematical model that *has* enjoyed a degree of success in the design of interactive systems is "Fitts' law", a "quantitative predictor" for movement time in targeting-type tasks derived by Paul Fitts in 1954 (Dennerlin et al., 2000) p423. Fitts' law, in its original formulation, is an empirical equation relating the time necessary to move one's hand between a starting point and a target area to 1) the target's width, and 2) the distance between the starting point and target. (For a mathematical formulation of Fitts' law, see Equation 1). Fitts used this model to

---

[12] Pun is Eric Gunther's. See Gunther et al., (2002).

describe three different targeting tasks (tapping plates with a stylus, placing washers onto pegs and placing pegs into holes) with the overall objective of building up an information-theoretical account of the human motor system, an account that could quantify human motions in terms of bits (Fitts, 1954).

$$MT = a + b \log_2 \left\{ \frac{A}{W_T} + c \right\}$$

$$ID_T = \log_2 \left\{ \frac{A}{W_T} \right\}$$

Equation 1: Fitts' law: an empirically derived relationship describing movement time *MT* as a function of movement amplitude *A* and target width *WT* along a direction of movement, where *a*, *b* and *c* are empirically derived constants. The logarithmic portion of the equation defines the *Index of Difficulty* for a particular targeting task (Dennerlin, 2000).

While Fitts' law was originally conceived to describe *physical* targeting activities, it has since been embraced within HCI research as a way to describe pointing and selecting in *virtual* spaces. Stuart Card first applied Fitts' law to the study and refinement of input devices for a computer, an application that contributed to the commercial introduction of the mouse (Card, 2002). Since this initial adaptation, numerous researchers have relied on (and sought to extend) Fitts' law as a practical tool for comparing the performance of alternate devices for pointing and selecting (Dennerlin et al., 2000) (Akamatsu and MacKenzie, 1996; Hasser et al., 1998). In the context of research in human-computer interaction, Fitts' law is typically applied from a Virtual orientation; it traditionally informs the design of *input* devices that allow people to feel as though they are *in* computational environments.

2.6. Environmental Concerns and Computation

As computers of various forms become increasingly widespread, their environmental impact increasingly warrants concern. This section briefly discusses two types of pollution that accompany digital technology; one manifest in the built environment, the other affecting the natural habitat.

2.6.1. Environmental Concerns of Computation in the Built Environment

While computational technologies imbue our houses, transportation systems, workplaces and commercial centers with wondrous capability, they also steep these spaces in unprecedented disturbance, distraction and disruption. Everywhere we turn, our senses are subject to bombardment by dynamic displays. The ring tones of mobile phones, the alarms of personal digital assistants, the buzz of pagers, email notifications, web site pop-ups, blinking electronic ads and billboards, application update reminders, the flicker of TV, radio jingles and the ambient industry of automated household appliances all vie for our attention and request engagement at one degree or another. In *Future Shock*, Alvin Toefler calls this problem "information overload"; in *Understanding Media*, Marshal McLuhan foretold an associated "age of anxiety" (Toffler, 1970; McLuhan, 1964). A current project at London's Royal College of Art advises the layman on the construction of "Digital [Emergency] Shelters": makeshift sanctuaries to stay the surrounding clamor of electronically mediated agendas (Sepulveda, 2004).

The barrage of electronic signals that pepper our built habitats can be viewed and described in environmental terms. John Thakara of the Netherlands Design Institute describes it as a pernicious form of "semiotic [symbol and sign-related] pollution"; Ezio Manzini of Milan's Domus Academy goes further, exorting designers to stop polluting the "semiosphere"(Thakara, 2003; Manzini, 1991). In the design of interactive systems, Ludwig Mies van der Rohe's dictum "less is more" begins to carry new ecological connotations.

Just why such distraction accompanies the presence of computational media is a subject of some debate. From an evolutionary perspective, the ability to attend and respond rapidly to sudden change is advantageous for survival, since danger often accompanies such change. Perceptual studies reveal that our eyes involuntarily saccade in response to "motion transients" in a visual scene (Intille, 2002). In addition to evolutionary and physiological explanations, there are pedagogical and commercial ones. Thakara notes that design cultures have historically been "obsessed with spectacle" and points out that designers are now trained to design messages, not interactions (Thakara, 2003). Stephen Intille of MIT's Home of the Future Consortium frames the problem of digital disruption as one of convergence and competition:

> The trend toward attention grabbing information clutter is a byproduct of the fact that…interfaces and…environments are not created based on a single, coherent vision of how information should be best conveyed to people…[thus] user experience emerges from *competing interests* (Intille, 2002) p93.

The blinking light (or sounding beep) of one electronic device may be innocuous, however the transient constellations produced by tens, hundred or thousands of displays predicted by proponents of Ubiquitous Computing threaten to leave us in a state of perceptual paralysis, as though dazed on the doorstep by the flash of paparazzi cameras.

Additionally, there are social and linguistic reasons why digital media so disruptively commands our attention. As Terry Winograd argues on philosophical grounds, computers are tools for language, and commitment is central to language acts (Winograd and Flores, 1986). Neil Gershenfeld more colloquially observes (in a discussion regarding email overload), "my mother taught me to speak when spoken to" (Gershenfeld, 1999) p104. We interpret transitions in digital display as human attempts to communicate, and such communication implicitly requests response.

Regardless of evolutionary, physiological, pedagogical, commercial, social or linguistic origins, the pollution of digital distraction within the built environment is a problem of considerable contemporary relevance. Mark Weiser at PARC described it as *the* "fundamental challenge for all technological design of the next fifty years" (Weiser and Brown, 1997) p76.

2.6.2. Environmental Concerns of Computation in the Natural Environment

Computation not only affects built habitats, it impacts the biological world as well; despite their clean image, computers hold increasingly dirty implications for the natural environment. Electronic "e-waste", most of which includes computation in some form, is one of the fastest growing waste streams in the industrialized world today. According to the U.S. Environmental Protection Agency, it now accounts for 1% of the United States' 210 million tons of annual solid waste, and this figure is growing (Rosenblith, 2004).

Electronic waste is particularly worrying because of the numerous toxic chemicals it contains. CRT Displays, for instance, contain on average of 4-8 pounds of lead, a heavy metal linked to human behavioral problems, learning disabilities, seizures and death (Rosenblith, 2004; EPA, 2003). The Polychlorinated biphenyls (PCBs) used to make circuit boards can cause a variety of adverse effects including "cancer, immune suppression, reproductive damage, birth defects, and fetal death" and tend to move up the food chain, accumulating in high concentrations within the bodies of larger predators, including humans (Temiskaming, 2002). Incineration of PCBs releases Dioxin, a toxin linked with various human and animal ills, including cancer, birth defects and diabetes (NIEHS, 2001).

Power sources for computational devices pose particularly serious environmental risks. Batteries are rich sources of heavy metals (such as Lead, Cadmium, Mercury) that accumulate in the bodies of animals and humans, with poisonous effects. This danger is particularly disconcerting given the current trends towards Mobile and Ubiquitous devices, and the marketing of computers as consumer items. These trends indirectly invite the disposal of more and more batteries in a less and less regulated fashion.

The extreme toxicity of the heavy metals used in batteries and the PCBs used in circuit boards – together with current trends towards Mobile and Ubiquitous devices – make it prudent to explore those exceptional computational technologies that do not inherently concentrate such materials, and to engage in such exploration at earlier and earlier stages in the design of computational products. (Environmental impact was one of the factors that was ultimately considered in choosing a tracking and identification technology. This particular design decision is discussed in Chapter 3, Section 1.4).

Whether operational in the built environment, or disposed of within the natural world, computers are a significant source of environmental pollution. For computers to become sustainable partners in human living and working, environmental issues of disruption and toxicity must be addressed more thoroughly than they are at present.

The past chapter has provided background on the topics that must be understood for the design decisions of this thesis to make sense. The following chapter draws from these disparate roots to define a research trajectory.

3. Path to Research Topic

This chapter charts my path towards an addressable topic of research. It begins with a sequence of key decisions to orient the reader, and ends with a research question. This question, in its rough form, motivated the preliminary explorations discussed in the next chapter, and ultimately led to the tighter, more addressable formulation introduced in the following chapter.

3.1. Early Decisions

Honing in upon a topic of research was a gradual intuitive process, however it is presented here as a series of discrete decisions for simplicity's sake. Decisions are summarized in the following way: Issues appear in bold, options are bulleted and chosen options are underlined. Rationales for choices follow the options.

3.1.1. Overall Approach to Human Computer Interaction

- Virtual: Strive to make people at home in essentially digital domains
- Embodied: Strive to coax computational advantages out into the human physical world

I chose to pursue the Embodied perspective because of reservations I have about the Virtual perspective. (These perspectives have been outlined in Section 2.1). I prefer to think of people as living and working in the physical world, with access to informational capabilities rather than vice-versa. I am concerned that the Virtual approach desensitizes people toward their own bodies' sensitivities and skillful capacities, and encourages a complacent trust in models and idealizations.

3.1.2. Embodied Approach to Human Computer Interaction

- Embedded Computing
- Augmented Reality
- Ubiquitous Computing
- Tangible Computing

Of the various Embodied perspectives, I was particularly interested in Augmented Reality for several reasons. Firstly, I was curious: could informational overlay actually enable best-of-both-worlds functionality? Would it do so without undermining established and trusted physical identities? Secondly, Augmented Reality demands clear added value from computation, and so I hoped this perspective would push my work beyond simple technological fascination. Thirdly, my background in engineering and classical animation made AR a natural perspective; the technical challenges of AR require engineering knowledge, and the artistic challenge of enhancing physical identity through the design of interactive behavior is, to some degree, analogous to the animator's challenge of enhancing character identity through the design of visual behavior.

3.1.3. Augmented Activity

- Finding Physical Things (I didn't consider a discrete set of alternatives)

I chose to focus on finding things, for a combination of historical and aesthetic reasons.

Historically, the ability to find (digital) things quickly and easily is a time-worn advantage of the digital domain. Algorithms for searching (together with sorting) have been a central theme in Computer Science (Cormen et al., 1990). Queriable databases drove the era of mainframe computing. Personal computing was defined largely through Apple's "Finder", a feature that remains with the Macintosh despite radical changes in software and hardware architecture over the past twenty years. In networked computing, search has emerged as a "killer application" in its own right. If the satisfaction people experience in ringing their own phone numbers to find misplaced cell phones provides any indication, support for finding things will be valued in situations where the computational "things" are significantly physical as well as digital[13].

---

[13] One of the original motivating scenarios for Ubiquitous Computing included in *The Computer for the 21st Century* concerns finding a lost (physical) instruction manual through digital means (Weiser, 1991).

Aesthetically, computation can be viewed as a medium and showcasing a medium's strengths is one way to create beautiful things. (This guiding aesthetic is encapsulated by Brancusi's maxim: "truth to materials") (Wilson, 2000). Although computation is not a material (and presents interaction designers with more even more plastic possibilities than plastic in industrial design) it *can*, like a material medium, be applied more and less appropriately. One thing computers do well is discriminate between discrete representations; this is essentially the digital strength. Using computation to support (comparison-based) search over a set of discrete (physical) items showcases an intrinsic strength of digital media in a new (physical) setting.

3.1.4. Means by which to Render Physical Things Computationally Findable (Means for Creating & Maintaining Physical-Digital Correspondences)

- Bar Codes
- Radio Frequency Identification
- Infrared Markers
- Ultrasonic Markers
- Magnetic Markers
- Colored Visual Markers

Given the decision to explore computational support for finding physical things, choosing a sensing technique assumed vital importance. I considered a variety of labeled techniques for maintaining physical/digital correspondence (unlabeled techniques were dismissed as infeasible), and ultimately chose Radio Frequency Identification based on its relative versatility, personal and local interest in RFID technology, and environmental considerations.

The properties of Radio Frequency Identification (RFID) make it a versatile tracking solution. RFID tags are wireless and battery-less. They don't require line-of sight or physical contact for identification. Tags can be remotely identified with a frequency approaching real-time. Thousands of tags can be identified and distinguished from one another. Multiple tags can be identified when all are simultaneously in range. RFID tags are cheap, and can be stuck to or embedded within a wide variety of physical items without significantly impacting their shapes. No alternative tracking modality combines such an unrestrictive set of advantages.

Personal and local interest also influenced the choice to focus on the potentials of the RFID tracking modality. The way RFID wirelessly modulates a power signal to communicate information struck me as an elegant way to simultaneously address multiple non-trivial technical problems, and I was curious about how well this would actually work. While I was reading about RFID's theory of operation, colleagues were exploring its practical application in the context of an interactive museum installation. I observed that RFID-tagged objects, used as control elements, could be somewhat confusing and frustrating; since the tags coupled with the interrogator through imperceptible means, creating favorable orientations between oneself, tagged items and the interrogator was a hit or miss pursuit. The sudden, discrete feedback provided by the interrogator seemed to compound the problem. I wondered: How would it feel if RFID provided something akin to a focal gradient to guide alignment? How might this change the quality of interaction? How might it change the frequency of successful interactions? While I began to consider this wide-open design space, students in the department's interactive media program began to ask about RFID with increasing frequency. In addition to being a versatile tracking modality, RFID was fast becoming a topic of personal and local interest.

RFID was also an attractive solution in light of its environmental impacts. Since RFID's mode of operation is not directly perceptible, RFID tags do not intrinsically add to busy-ness of signs and signals within the built environment. Since RFID tags require no batteries and are not composed of PCBs, their impact on the natural environment is far less drastic than the impact of other electronic label-based tracking technologies[14]. Environmental issues may seem to be a world away from issues of interface and interaction, however designing responsibly from an Embodied perspective (i.e. computation everywhere) requires that they be considered.

---

[14] Additionally, RFID presents new opportunities for product life-cycle monitoring and automated sorting, which may have implications for recycling practices. Such implications are, however, intertwined with notions of personal versus public property, and thus privacy.

3.1.5. Focus within Finding Physical Things with Label-Based Digital Assistance

- Tagging & categorizing items so that they might be found with digital assistance
- Specifying what one is looking for to the supporting system
- <u>Locating a sought item that has been specified (an item that has already been tagged and included in the set of things that can be sought with digital assistance).</u>

Enabling automated support for finding digitally labeled physical items is a challenge with many parts. Tagging, categorizing, specifying and locating must all be addressed well before a system might provide useful digital support for finding physical things.

I chose to focus on locating (in the sense of getting one's hands on) a previously labeled and cataloged item, once it has been specified as a sought item. This choice was made based upon personal interest in the technical and interaction-related issues involved. I was curious about RFID's potential to provide continuous distance estimates, and interested in considering what sort of feedback a system for assisting location might provide to its users.

3.1.6. Scale of Locating Assistance

- City
- Building
- Room
- <u>Shelf</u>

The maximum range of RFID systems is presently limited to a few meters. With the aim of keeping this research grounded with respect to what is possible rather than what is notional, I chose to focus on the scale of shelf-sized spaces. Shelf-sized spaces of roughly arms-length proportions appear throughout home and office environments of the western world; if something is lost, it often can be isolated to one or another container (or delimited region) of shelf-sized proportions.

3.1.7. Spatial References for Automated Assistance

- Absolute Coordinate System, embedded in environment
- Relative, from hand tool; Reports Proximity Only
- Relative, from hand tool; Reports Direction Only
- Relative, from hand tool; Reports Direction & Proximity

Digital assistance for locating physical things might be provided relatively or absolutely. (Shining a light on a sought item to indicate its location in a room is an example of absolute assistance; the seeker's position relative to room or item is not a factor influencing the assistance given. Pulling one's hand magnetically towards a sought item, on the other hand, is an example of relative assistance; the seeker's location influences the assistance given.) Relative assistance can communicate proximity, direction, or both. A Geiger counter (for example) reports relative proximity; only through perceiving its movements over time can one obtain a sense of direction and steer towards source of radiation. Dowsing rods that appear to orient themselves towards water report relative direction[15]. Relative assistance does not have to assume hand-held form, however I chose to consider only hand-held solutions. Hand-held implements can be taken up and put down as needed, and thus provide auxiliary rather than enforced support (this is in line with the AR approach of Augmentation rather than replacement). Furthermore, shelf-sized spaces are easily traversed by hand. Finally, if a sought item is to be retrieved by hand, it is ultimately the hand that needs to be guided to it.

I chose to focus on relative solutions rather than absolute solutions based on the belief that digital assistance for finding things should not disrupt existing physical systems of categorization – once more, based on the aim to augment rather than replace. People are pilers as well as filers; we place things idiosyncratically as well as in (objectively) discernable orders. Absolute solutions limited by shelf-scale maximum range are likely to miss (for example) the pile left by the door on the way out, the articles left on the chair, the things that fell in the space behind the desk, and

---

[15] In dowsing, it is actually supposed that the person guides the stick (Power, 2002).

numerous other exceptional placements. A relative solution, in contrast, can leverage the spatial knowledge and experience of its wielder.

In consideration of feasibility, I chose to focus on communicating proximity rather than direction (or a combination of the two). Sensing direction through RFID is a more difficult technical challenge than sensing proximity; it entails a greater number of circuit elements and ultimately, a greater cost. Furthermore, communicating direction to a person through a hand-held device is also more complicated than communicating proximity. Delivering a varying stimuli over the surface of a hand-held device or making a device (or part of the device) physically twist and point would be considerable engineering challenges in their own right.

Given these decisions, automated assistance for locating things can be reduced to communicating the relative proximity between a sought item and a seeker's hand.

3.1.8. Continuous Guiding Stimulus

- Light
- Sound
- Force
- <u>(Tactile) Vibration</u>
- Multimodal Combination (of two or more of the above)

While light and sound provide clear signals and are perceived with extensive dynamic range, they are by nature *public* stimuli that can erode calm in environments shared by multiple people and devices. If, as the AR worldview maintains, the physical world is to be sprinkled extensively and beneficially with computational capability, private channels of communication are worth considering in the interest of minimizing disruption.

Haptic stimuli are by nature private, and thus introduce less potential for distraction than visual or auditory stimuli. Additionally (as mentioned in the context of the previous decision) if a sought item is ultimately to be retrieved *by* hand, it seems natural to provide guiding stimuli *to* the hand. Force-feedback is relatively expensive, heavy and power-intensive, and so I chose to focus on vibrotactile feedback as a means to communicate the proximity of a sought item to one's device-holding hand.

3.1.9. Modulation of Guiding Stimulus (to Communicate Distance)

- Amplitude Modulation (AM)
- Frequency Modulation (FM)
- Composite AM/FM

A vibrotactile stimulus might be modulated (varied) in many ways to communicate "close", "far" and the spectrum in between. The three forms of modulation I considered were Amplitude Modulation (AM), Frequency Modulation (FM) and the combination of AM and FM[16]. Since combined AM/FM can result in wider perceived dynamic range than either AM or FM alone (Murray et al., 1998), this approach seemed, at first, to be ideal. Unfortunately, the only vibrotactile transducer I could find that was a) robust and b) small enough to fit within a hand-held device lacked the flat frequency response necessary for controlled delivery of a composite AM/FM stimulus. Rather than attempt to introduce compensation for non-flat frequency response, I decided to simply use AM to communicate "getting closer" or "getting farther"). Amplitude Modulation was far simpler to implement and verify, so I was practically limited to exploring the modulation of amplitude alone.

3.1.10. Target Approach to Evaluation

- Create a working system for one specific application, let people use it, see what they think.
- Conduct quantitative laboratory experiments of the human factors tradition, based on simulated functionality.

While narrowing down a topic of exploration, I began to consider approaches for final analysis. The two approaches I considered were: 1) create and refine a

---

[16] Conceivably, modulation of a sense-able parameter to convey "closer" or "farther" could be dynamic – it could change over time. For example, in the childrens' game "find the spoon", the point at which "cold" changes to "warm" might grow closer to the spoon during the course of play. In this thesis, only static mappings between distance and the guiding stimulus are considered.

prototype for one specific application 2) conduct quantitative laboratory experiments of the human factors tradition based on simulated functionality in the context of a generalized task.

I chose the second approach based on academic considerations and feasibility concerns. While I could envision arriving at falsifiable generalizable results, analyzing with rigor and delineating a scientific contribution through the second approach (all expectations of a Master of Science program), I could not envision meeting these expectations through the first approach.

Prototyping physical and electromagnetic/RF devices can be time consuming, and arriving at even one robust functional prototype for a handheld locator within a two-year timeframe would be an ambitious goal. If such a prototype were constructed, it would be just one small part of a much larger system. The rest of the system (the mechanisms for tagging, cataloging & specifying sought items discussed in decision 4.1.5) would need to be implemented, and the complete system then matched appropriately to an actual finding task before any truly realistic evaluation might take place “in the wild” (Gaver, 2002). Even if all of these preconditions for a situated, realistic evaluation were satisfied over the course of a two-year time frame, evaluation could easily require additional time; a system for locating things might show its true value over years of use, rather than over minutes or days.

Ultimately, moving an idea part of the way out into the world in the form of a useful product (the spirit of the first approach) would be hugely satisfying, however I chose to pursue the second approach because it seemed more appropriate for the constraints of an MSc. program. Additionally, physical search and targeting tasks have historically proven to be somewhat amenable to quantitative empirical study.

3.2. Research Topic (Rough Cut)

Given these decisions, my research question now became:

How can the amplitude of a vibrotactile stimulus delivered to a person's hand through a hand-held device best be varied so as to guide the hand towards a sought item in a shelf-sized space, assuming the sought item has been specified, labeled with an RFID tag, and catalogued within a set of "findable" items?

Though this first formulation was still somewhat loose, it proved sufficient for guiding preliminary explorations (discussed in the next chapter) towards a tighter, more answerable formulation.

4. Preliminary Explorations

A hand-held locator of labeled items is, at this point, largely a hypothetical device; in order to design and evaluate interactive feedback for one, it is first necessary to conceptualize what such a locator might be like and how it might work. Next, it is necessary to arrive at a prototype or simulation that is sufficiently *like* an actual hand-held locator so as to support the delivery and evaluation of any interactive behaviors that are designed. Toward these ends, several preliminary explorations were conducted. These explorations served to ensure the feasibility of a proposed underlying implementation, tighten the topic of research (i.e. clarify which aspects of interaction to focus on in depth), create a stable platform for experimentation, and ultimately, to inform the design of evaluative experiments. The preliminary exercises began with scenarios, continued with technical exploration and simulation, and culminated with a functional demonstration.

4.1. Scenarios

4.1.1. Purpose

In order to move the somewhat abstract idea of finding things through a vibrotactile gradient into a public space for discussion and critique, fictional scenarios were written and shared to illustrate the concept at work. These scenarios situated the envisioned functionality within several human activities, and delineated its intended utility in each. Through sharing and refining scenarios, it became possible to move from private, personal intuition towards a more public, interpersonal intuition[17].

4.1.2. Scenarios

What, practically, would it mean to have automated assistance for locating everyday objects in shelf-sized spaces, assistance that is provided through a vibrotactile

---

[17] It is perhaps too much to claim that the sharing of scenarios verified anything beyond shared intuition, since activities and protagonists were constructed rather than "found" through actual user research or demographic study.

feedback gradient from a hand-held locator? The following scenarios illustrate how such gradient-based assistance might prove useful.

4.1.2.1. Easing Eye Contact

Rupa is a cashier at Wireless House, a small consumer electronics store. She rings up hundreds of products each week, and her handling of individual products during checkout provides customers with the opportunity to ask last minute questions. Rupa can "feel" the location of product identification labels on packages through the handle of her hand-held scanner. Since she can locate and scan Ids on the packages largely by touch, she is free to maintain eye contact and conversation with customers more easily during checkout than she could while using a traditional product ID scanner. Touch adds a nuance to her handling of the hand-held scanner, a nuance that facilitates hundreds of transactions each week.

4.1.2.2. Dowsing for Documents

Alfred, a research scientist, dreams of the ability to summon documents to his fingertips instantaneously from the cluttered corners of his office. Locomotive magic pending, he settles for automated assistance[18]*. A wand with twitching behavior akin to a dowsing rod's facilitates locating papers for retrieval, while not precluding other techniques for locating nor interfering with his existing systems of organization that suffice most of the time.*

4.1.2.3. Supporting a Soundscape

Nathan, a sound engineer and environmental activist, wishes to juxtapose urban and pristine wilderness sound recordings in an exhibit calling for ecological responsibility. The exhibit consists of a large globe centrally located in a room and a rack of stethoscopes to one side. Upon entering the room, visitors receive a stethoscope and the injunction to examine the "pulse of the planet". By moving stethoscope heads over the globe surface visitors tune between spatially mapped sound recordings as if tuning between stations on an analog radio. Low frequencies are translated into vibrations of the stethoscope head; this vibration facilitates

---

[18] This scenario presumes a) that there is an effective mechanism for registering documents that enter Alfred's collection and b) that Alfred has an easy way to specify the document he seeks. Although these are considerations of great importance, they are outside of the scope of this project.

location of planetary sound sources, improves perceived sound quality and communicates the sense that the planet is alive. Visitors leave the exhibit having considered their own relation to the earth in terms of a doctor/patient relationship and having had the opportunity to traverse their own unique paths through a sonic world.

In each of these scenarios, the ability to sense physical proximity via virtual touch facilitates a locating activity. Tactile location cues allow Rupa to attend more fully to customers while finding product labels, provide Alfred with an auxiliary way to locate documents from his paper archives and offer Nathan an appropriate medium for his ecological message.

4.1.3. Reactions

These and other scenarios were informally shared with colleagues for their review. While various people embraced/rejected different scenarios to varying degrees, all reviewers deemed at least one scenario to be a compelling possible use for vibrotactile location cues. (Many of the people available for comment were colleagues and "early adopter" types with an interest in interactive systems, so a somewhat soft reception was anticipated). Numerous reviewers expressed the desire to just try it out: to experience what their favorite scenarios described. I interpreted this interest and eagerness as an indication that continued pursuit was worthwhile in local as well as personal terms.

I had hoped that the process of writing, sharing and discussing scenarios would suggest a particular application domain to keep in mind for subsequent user research and demographic inquiry, however this did not take place. The main outcome of working through a fictional, narrative mode of discourse was an understanding for how the idea of locating through vibrotactile feedback could and (and couldn't) be effectively communicated. Through discussing the scenarios, it soon became clear that words such as "search", "location", "identity", "label" and "feedback" are all ambiguous (they have various meanings in physical and virtual contexts), and that certain metaphors, such as the dowsing rod and Geiger counter, were particularly effective for clarification. The discovery of ambiguities and descriptive metaphors early on in the design process helped to facilitate communication throughout the project.

4.2. Technology Exploration

4.2.1. Purpose

In the interest of keeping research grounded with respect to what was technically possible[19], I began an in-depth, hands-on exploration of Radio Frequency Identification (RFID), the pivotal enabling technology[20]. By building and refining an RFID interrogator from published plans, I sought firstly to become aware of the main technical hurdles[21] and secondly to demystify the technology. RFID is presently enshrouded by commercial hype, so it seemed worthwhile to cultivate direct and personal experience regarding its actual potentials.

In addition to fostering a practical understanding of RFID's main technical issues and actual capabilities, there was the possibility that development might ultimately yield an apparatus for experimenting directly with mappings between tag signal strength (distance) and the intensity of a vibrotactile stimulus.

4.2.2. Starting Point

The starting point chosen for technical exploration was a tag reader design published by Microchip Technology to promote their 13.56MHz tag transponder chips. This design was chosen for a number of reasons. First, it was accompanied by thorough documentation. Second, it was a current design (Microchip's own RFID evaluation kits were based upon it), so knowledgeable support was available. Third, since the

---

[19] In considering the matter of implementation, interface designers often adopt the stance that "reality bats last" (Cooper, 1999). While this perspective is useful under certain circumstances, the circumstances presented by experimental sensor systems are far from certain. Unlike processors and memory, sensors do not, in general, grow faster smaller and cheaper year after year. Their behaviors can be quirky, and their limitations subtle. If an interaction design depends on experimental sensor systems, sensing must be considered up front in a design cycle for this cycle to produce a certainly realizable result.

[20] Vibrotactile transducers required no such extensive hands-on investigation; their mode of operation is simple and familiar, and the chances of designing an ultimately unrealizeable vibrotactile stimulus seemed low.

[21] RFID combines radio-frequency and power electronics, two areas full of surprises for the uninitiated.

design was for a reader (rather than a bidirectional interrogator), reader/tag communication protocols were relatively simple to implement and verify. Finally, the design's resonant frequency of 13.56MHz was suitable. This frequency is currently licensed for RFID applications in the E.U. and U.S. (and for Industrial/Scientific/Medical purposes world-wide)[22], it is close to the frequency that maximizes read range for inductively coupled systems (~10MHz) and it is not so high as to make manual prototyping impractical (Finkenzeller, 2002).

4.2.3. Development Process

The reference design consisted of four subsections: a transmitter for sending power to the tag, a receiver for receiving tag signals, a controller for interpreting tag signals and communicating with a host PC, and a power supply. (A high-level block diagram depicting these subsections appears in Figure 13 below).

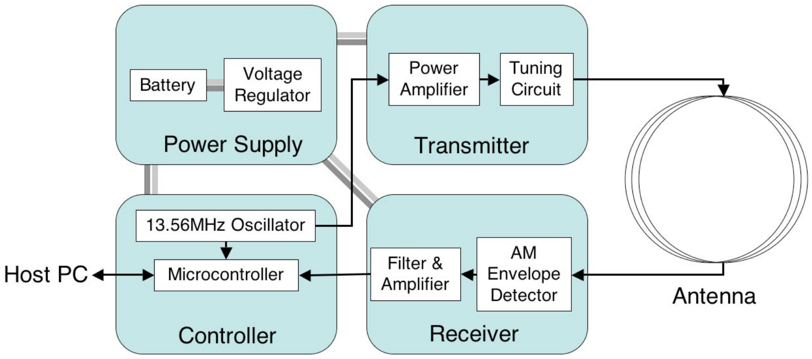

Figure 13: High-level block diagram of an RFID reader showing power, control, transmit and receive subsections. A 13.56MHz signal is amplified and transmitted through the same antenna used to receive incoming tag signals. The receiver detects, filters and amplifies the amplitude-modulated "envelope" of an incoming tag signal. It then converts the signal to a series of digital transitions for the control subsection to decode and relay to a host computer system.

Functional blocks were implemented and refined independently where possible, and together where necessary. (Receiver and transmitter shared a common antenna, thus these two functional blocks were highly interdependent.) After the various subsections had been created, tested, joined and tuned through traditional hardware and firmware techniques, the resulting system could distinguish between two "test" tags,

---

[22] Theoretical calculations regarding the legality and range of a hand-held RFID system are presented in Appendix E.

albeit in brittle fashion. Photographs of the realized system appear in Figure 14 and Figure 15 below. Schematics for the initial and refined designs appear in Appendices A and B. A listing of the firmware for the control subsection appears in Appendix C.

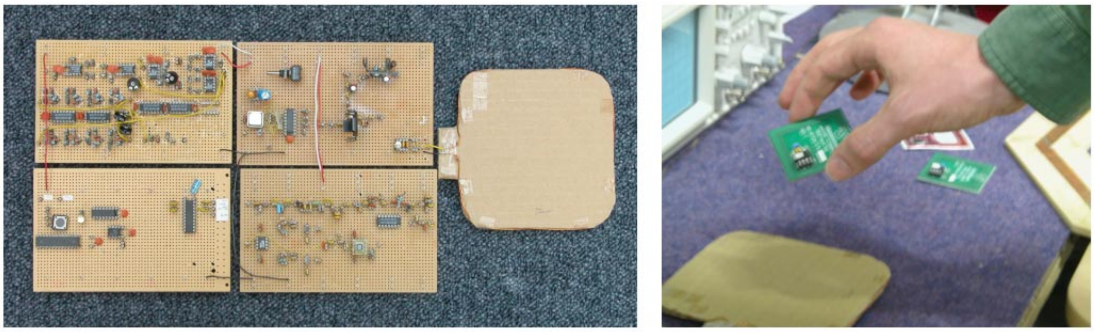

Figure 14: Left: The implemented RFID tag reader. Right: Moving a test tag above reader antenna.

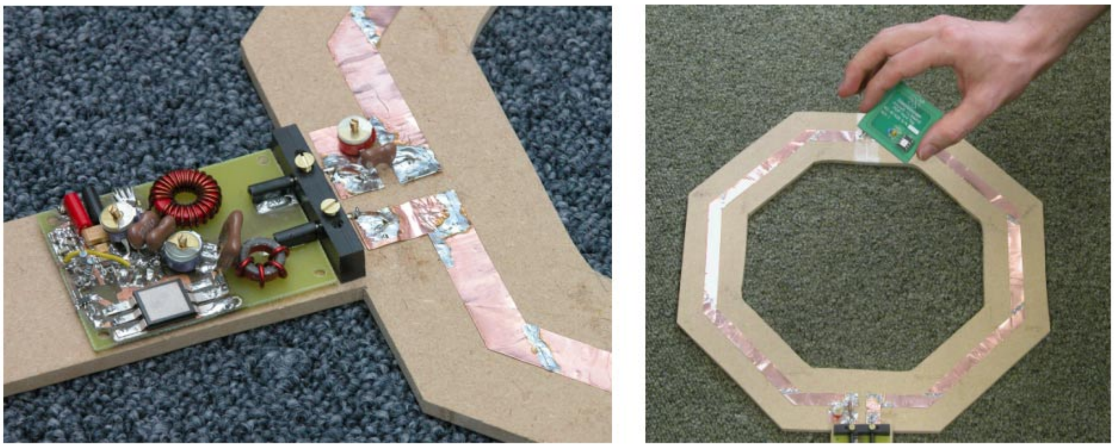

Figure 15: Left: Revised transmitter subsection. Right: Moving a test tag above an over-sized antenna. This antenna successfully powered the resonation and switching of RFID tags at a distance of approximately 1.2 meters.

4.2.4. Lessons Learned

The process of constructing and testing an actual RFID system yielded several practical lessons. Some were relevant to my focus of designing a guiding stimulus for a hand-held locator, while others were more purely concerns of implementation. Lessons relevant to designing a guiding stimulus are presented below, while lessons concerning implementation are included in Appendix D.

One major concern I had upon commencing technical exploration was whether or not a maximum range of one meter was actually feasible for a hand-held reader and credit-card sized tags. Since inductive power transmission between (axially aligned) antenna coils decreases rapidly with distance (power is proportional to the inverse sixth power of the distance between a reader and a tag) it initially seemed that powering tags at one meter's distance would prove to be the toughest practical problem (Sorrells, 2000). Testing of the transmitter and receiver sub-sections revealed that a range of a meter *was* feasible; in my implementation, isolating tag transmissions from noise and the powerful 13.56mHz carrier proved to be a much more difficult challenge. (Note: Since this investigation, hand-held RFID readers of a two-meter maximum range have become commercially available.)[23]

Another initial concern was the actual rate with which tag IDs could be reported; if the report rate proved too low, the illusion of real-time continuous feedback might not be sustainable. The datasheets for Microchip's MCRF355 tag transponder chips suggested a theoretical maximum frequency of 10 tag reports per second (Lee and Sorrells, 2001). Was this frequency practically realizable? If so, would it prove sufficient for guiding hand movement in "real time"? I attempted to test these concerns in an extremely limited manner by moving a tag held at arms length above the stationary reader antenna and watching ID reports scroll by on a terminal window. When tags were in range, reports did in fact occur ten times a second; near maximum range, reports became intermittent. Since no measure of signal strength was actually reported by the reader (only ID numbers were reported), I had to imagine that a value was being updated with each new ID report, a value that could be translated to a stimulus for guiding hand position. Within this extremely crude, subjective and imaginative test, the report rate seemed to support the illusion of continuous control for movements of the arm, but seemed not to do so for quicker movements of the wrist and hand. Given that searching a shelf-sized space with a hand-held device would likely require extensive arm movement – and that the weight of a handheld battery-powered device would filter high-frequency hand motions – a report rate of ten times per second might be sufficient to support the illusion of continuous control, at least within the context of a rough demo. (A higher rate would

---

[23] For a theoretical treatment of read range and legality issues, see the calculations presented in Appendix F.

be preferable, and data sheets for other commercially available tags suggest that higher rates are possible) (Texas Instruments, 2002). The impact multiple in-range tags had on each other's report rate (and signal strength, as measured by the reader's antenna) was not investigated, though it should have been.

4.2.5. Decision to Simulate

After allaying feasibility concerns and identifying main implementation issues through technical exploration, I was faced with a decision: Continue development of an experimental apparatus in hardware – a device capable of actually mapping tag signal strength to the intensity of a vibrotactile stimulus – or simulate such a mapping, and base further experimentation on this simulation.

I chose to proceed with simulation, because my experience of building and verifying an RFID reader suggested that development of an experimental apparatus in hardware would not be possible within my timeframe. The comparatively short development time required for simulation would allow me to focus on the main topic – an interactive behavior for an appliance – without becoming sidetracked by issues of implementation. (Since I'd just explored what these technical issues were, however, it would remain possible to keep the simulation *accountable* to an end implementation.)

4.3. Simulation of Functionality

4.3.1. Starting point

While scenarios made vibrotactile cues for locating accessible to the mind's eye, and technical exploration suggested their feasibility, neither of these preliminary investigations provided a means to experience what it actually feels like to locate by hand through a vibrotactile gradient. To support this perceptual experience, a simulation was constructed. The simulation was composed of a real-time 3D positional sensor (to assess distance between one's hand and a sought item), a software application (to map hand-item distance to the amplitude of a vibrotactile stimulus), and a vibrotactile transducer to convey the vibrotactile stimulus to the "seeking" hand. (For a diagram depicting these elements, see Figure 16).

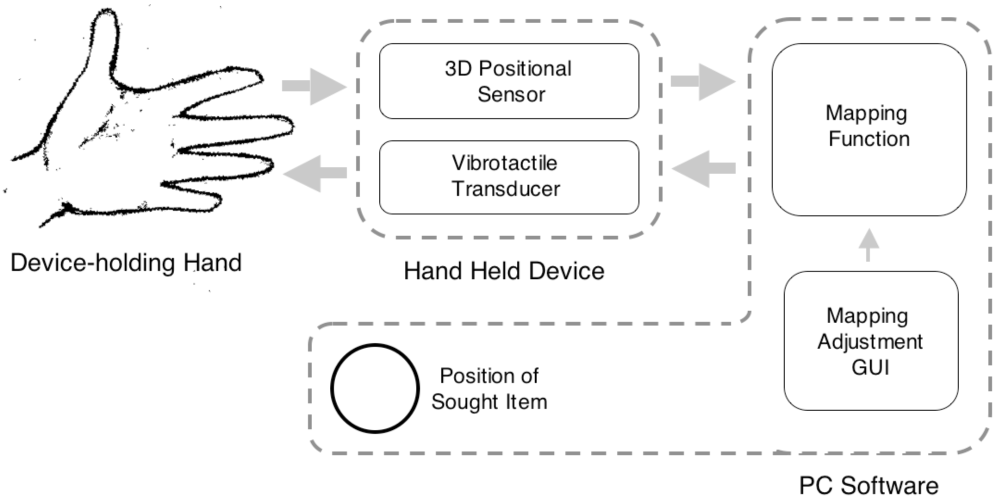

Figure 16: High-level block diagram of the apparatus used for simulating vibrotactile location cues. A sensor embedded in a hand-held device conveys the position of the hand-held device to a software application. This software application calculates the distance between the hand-held device and the position of the sought item, then maps this distance to the amplitude of a sinusoidal signal. The sinusoidal signal is sent to the hand through a vibrotactile transducer, and can ultimately be used to guide movement.

4.3.2. Development Process

The main areas of simulation development concerned delivering the vibrotactile stimulus, keeping track of the positions of the locator and sought item and mapping the distance between locator and sought item to the amplitude of the vibrotactile stimulus.

The sinusoidal waveform used to drive vibrotactile display was synthesized by a software signal generator (programmed in Snack[24]), and delivered through a PC's sound port to a VBW32[25] Skin Stimulator via a power amplifier. The VBW32 transducer was chosen because it was robust, efficient at vibrotactile frequencies (roughly 50-500Hz) and small enough to be incorporated within a hand-held device (Murray et al., 2003). The frequency of vibration was set at 250Hz, since the skin of the hand's digit pads is most sensitive to vibrations near this frequency (Erp, 2002).

---

[24] http://www.speech.kth.se/snack/

[25] http://www.tactaid.com/audiologicalengineering/skinstimulator.html

A picture of the vibrotactile transducer and power amplifier appear in Figure 17 below.

Figure 17: Left: The VBW32 transducer used for vibrotactile display. Right: A custom amplifier used to boost the power of the signal delivered to the transducer.

To sense the position of the locator, a Polhemus Fastrak[26] (see Figure 18) was used for its availability and simplicity. (Unfortunately, the available Fastrak could not report relative orientation, a limitation with implications for how well spatial aspects of tag/reader communication – and thus item/locator interaction – could be modeled. These implications are discussed in 4.5.3.).

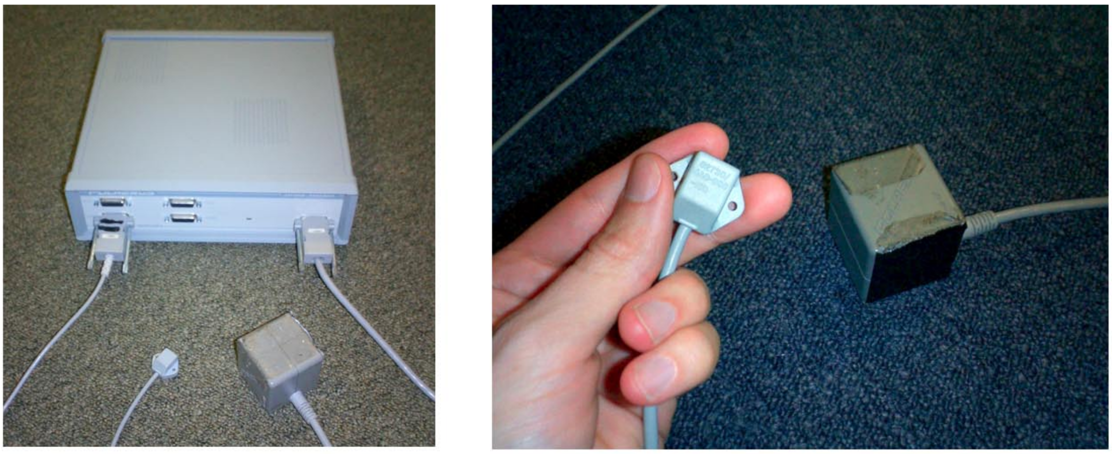

Figure 18: Left: The Polhemus Fastrak sensor system used to sense the position of the locator in three-dimensional physical space. Right: A close-up of a Fastrak sensor unit. In the background is the device that serves as the origin for the Fastrak's coordinate system.

---

[26] http://www.polhemus.com/

While locator positon was tracked by the Fastrak, sought item position was not. In this simulation, the sought item was a virtual construction; its location could be set automatically by software, or manually by moving the locator to a desired position and pressing a button. Making the sought item a purely virtual construction had two main advantages. First, it permitted arbitrary items at hand to be quickly and easily "set" as the sought item (an advantage for demonstration). Second, it permitted random and iterative placement of the virtual sought item's location (an advantage for experimentation).

Making the sought item a virtual construction also created a potential source of confusion: If a sought item location was set as a location in space where no physical item was visually or tangibly present, there would be no clear way to recognize when the locator had been successfully steered into it. To make such virtual placements directly perceptible, a pulsation was incorporated into the simulation, a pulsation between maximum amplitude and no feedback at all. This pulsation indicated when the locator occupied the space corresponding to the sought item (in software, the sought item was modeled as a sphere). The frequency of pulsation was set to 10Hz so that the resulting stimulus *in* the sought item's space would feel qualitatively different from the vibrotactile feedback used to guide the locator *towards* the sought item. Much time and effort could have been invested in the "in sought item" stimulus; I chose to implement it quickly and press on because this stimulus was essentially a requirement for the simulation, not a designed behavior for the actual implementation.

Since the Fastrak position sensing system was not always available and could be time consuming to set up, a mouse outfitted with a vibrotactile transducer was used as a provisional input/output device. A picture of this makeshift device appears in Figure 19.

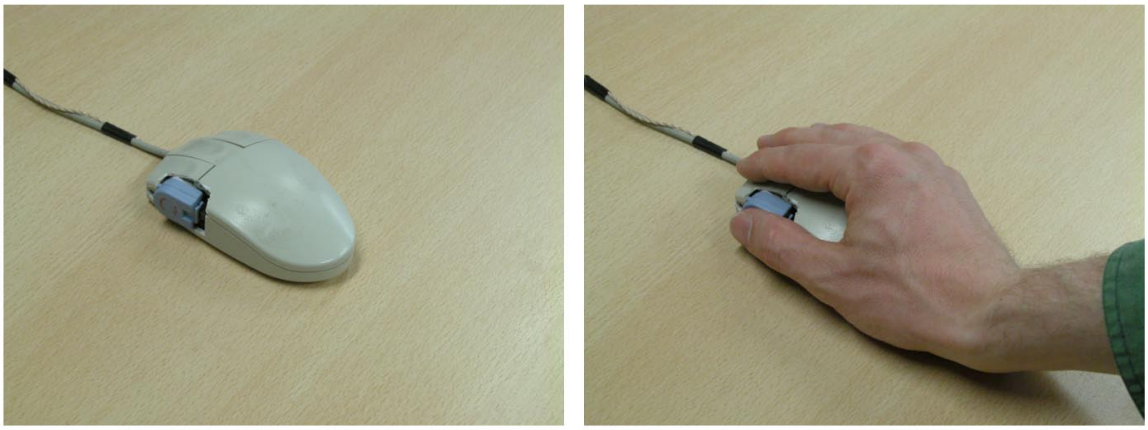

Figure 19: The mouse and vibrotactile transducer used provisionally for input/output during development of the simulation.

Obtaining positional information through the Fastrak (mouse) and delivering vibrotactile stimuli through the VBW32 were straightforward implementation tasks; the main challenge of developing the simulation was to create software for managing the function that mapped between locator-item distance and amplitude of vibrotactile feedback.

The software application used to map between distance and vibrotactile amplitude was a Tcl/Tk[27] script running on a 555GHz PC. To create this software, I began with a simple graphical representation of the correspondence between distance and amplitude, and progressively added ways to modify this correspondence. The aim was to include a variety of adjustable parameters within the simulation, so that the most relevant for locating might be identified later through experimental means. The simulation's adjustable parameters were:

<u>Locator Position</u>: Locator Position was the position of the hand-held locator (Fastrak sensor) within the search space, relative to the Fastrak's coordinate system. When a mouse was used as the locator, Locator Position was the position of the mouse pointer above a black rectangular "viewing screen" within the software application's graphical user interface (GUI). (The GUI appears in Figure 20, the viewing screen appears in the figure's lower left-hand quadrant).

---

[27] http://www.activestate.com/

Target Position: Target Position was a set of coordinates representing the position of the sought item in the search region, relative to the Fastrak's coordinate system. (Sought item locations were set in software rather than sensed by any sensing system.)

Input Source: The Simulation Input menu provided a way to switch back and forth between the mouse and Fastrak, between 2D and 3D search regions.

Target Size: Target Size provided a way to adjust the scale of the sought item. (The sought item was represented as a spherical region, and Target Size determined the radius of this spherical region).

Mapping Function: The Mapping Function menu provided a way to select between various predefined mathematical functions (step, ramp, exponential, etc.) for mapping between locator/item distance and amplitude of vibration. Mapping functions could also be "drawn" with a mouse directly upon the GUI's display of the current mapping function. (See Figure 20, upper left-hand quadrant).

Mapping Exponent: To make it possible to compensate for the psychophysical reality that perceived amplitude of vibration is not the same as actual amplitude of vibration, an exponent was added to the mapping function. This exponent was made adjustable to allow for experimentation.

Maximum Range: The Maximum Range was the distance at which the mapping function ceased to map between distance and amplitude of vibration.

Maximum Amplitude: Maximum Amplitude set the ceiling for the amplitude of the vibrotactile stimulus.

Distance Resolution: Distance Resolution determined the granularity with which distance was related to amplitude values by the mapping function.

Amplitude Resolution: Amplitude Resolution determined the granularity of the vibrotactile stimulus' amplitude.

These parameters were displayed and controlled through the graphical user interface shown in Figure 20. The interface's viewing screen also provided a way to indirectly check whether or not the simulation was functioning; if either the Fastrak or software application stopped working properly, the graphical representation of the locator on the viewing screen (see lower left-hand quadrant) would cease to reflect the actual movements of the locator.

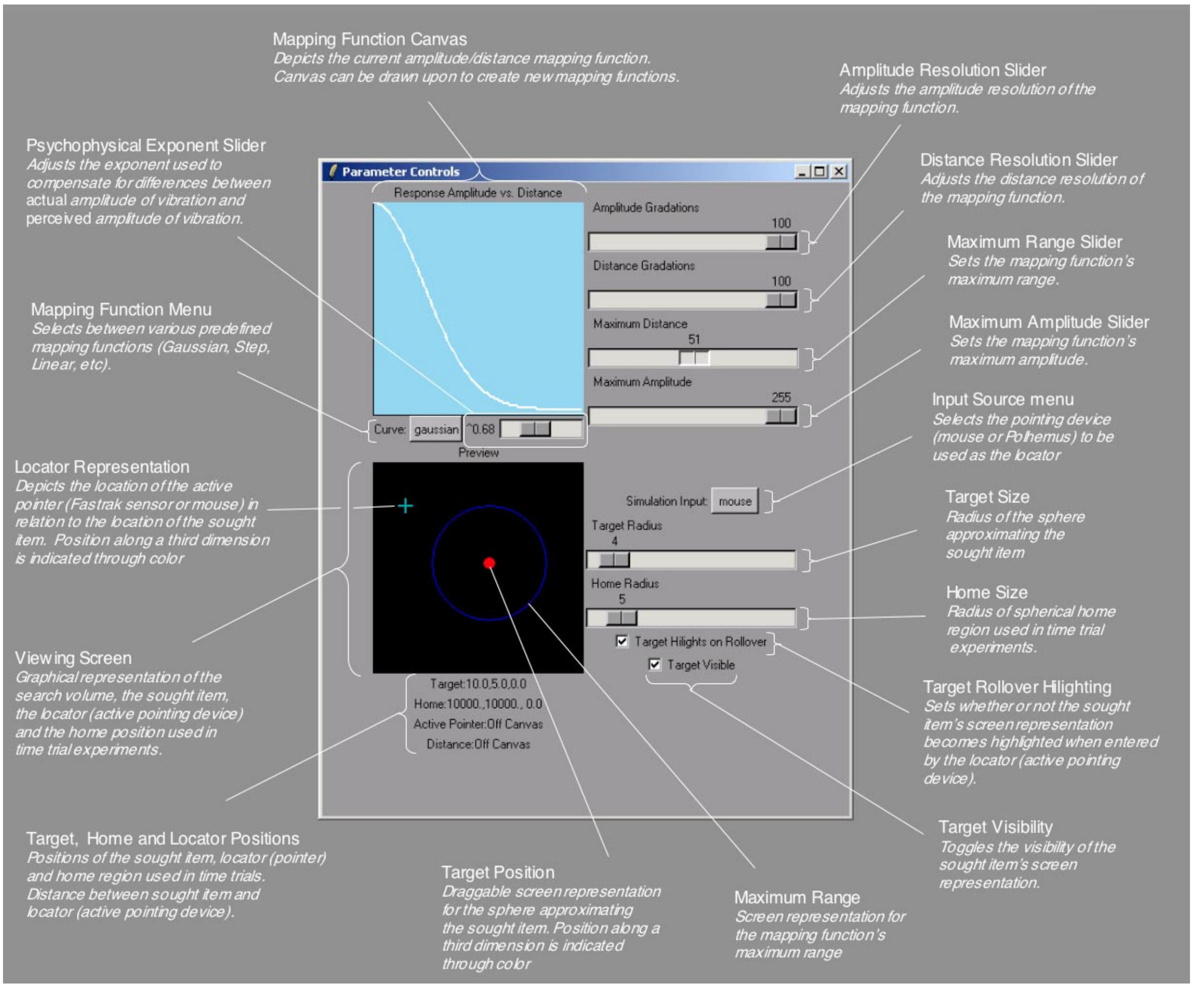

Amplitude Gradations 100

Distance Gradations 100

Maximum Distance 51

Maximum Amplitude 255

Curve: gaussian ^0.68

Simulation Input: mouse

Target Radius 4

Home Radius 5

☑ Target Hilights on Rollover

☑ Target Visible

Target:10.0,5.0,0.0

Home:10000.,10000., 0.0

Active Pointer:Off Canvas

Distance:Off Canvas

Mapping Function Canvas
Depicts the current amplitude/distance mapping function. Canvas can be drawn upon to create new mapping functions.

Amplitude Resolution Slider
Adjusts the amplitude resolution of the mapping function.

Psychophysical Exponent Slider
Adjusts the exponent used to compensate for differences between actual amplitude of vibration and perceived amplitude of vibration.

Distance Resolution Slider
Adjusts the distance resolution of the mapping function.

Maximum Range Slider
Sets the mapping function's maximum range.

Mapping Function Menu
Selects between various predefined mapping functions (Gaussian, Step, Linear, etc).

Maximum Amplitude Slider
Sets the mapping function's maximum amplitude.

Input Source menu
Selects the pointing device (mouse or Polhemus) to be used as the locator

Locator Representation
Depicts the location of the active pointer (Fastrak sensor or mouse) in relation to the location of the sought item. Position along a third dimension is indicated through color

Target Size
Radius of the sphere approximating the sought item

Home Size
Radius of spherical home region used in time trial experiments.

Viewing Screen
Graphical representation of the search volume, the sought item, the locator (active pointing device) and the home position used in time trial experiments.

Target Rollover Hilighting
Sets whether or not the sought item's screen representation becomes highlighted when entered by the locator (active pointing device).

Target Visibility
Toggles the visibility of the sought item's screen representation.

Target, Home and Locator Positions
Positions of the sought item, locator (pointer) and home region used in time trials. Distance between sought item and locator (active pointing device).

Target Position
Draggable screen representation for the sphere approximating the sought item. Position along a third dimension is indicated through color

Maximum Range
Screen representation for the mapping function's maximum range

Figure 20: The graphical user interface used to adjust simulation parameters.

Once spatial input, vibrotactile output and mapping management had all been implemented and coordinated, my focus turned toward the form of the hand-held locator: a factor that would likely have a significant impact on how people encountered it and made sense of its behavior. While the unpackaged Fastrak sensor element (and alternately, the mouse) had been sufficient to convey a sense of the locator's physical presence and location during development, they would not suffice for more polished demonstration and experimentation.

My limited exploration of form consisted of sketching, scavenging, constructing and sharing. I drew sketches to explore a range of possible forms and metaphors, scavenged products with forms that seemed in some way or another to convey the idea of a hand-held locator, constructed mockups from plastic and cardboard, and informally shared the results of these activities with colleagues for their reactions. Examples of scavenged and constructed forms appear in Figure 21. Sketches and doodles are included in Appendix G.

Figure 21: Shapes scavenged and constructed to explore possible forms for a hand-held locator.

The form ultimately chosen to use (see Figure 22) was selected as one that might simultaneously suggest functionality, support an underlying implementation and facilitate demonstrations and experiments.

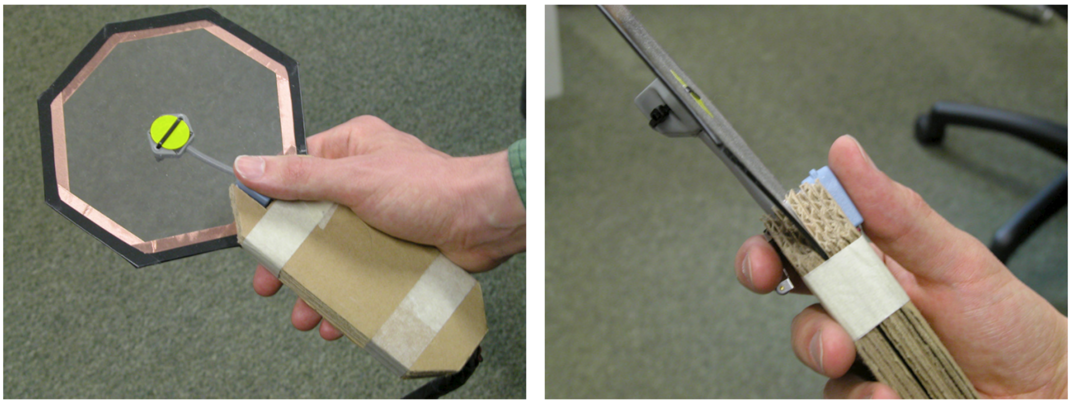

Figure 22: Left: The hand-held locator's final form. Right: A close-up illustrating the locator's active elements: a vibrotactile transducer (under thumb), a button (activated by one's trigger finger) and the Fastrak position sensor (visible to the upper-left).

The locator's form was abstracted from a magnifying glass, a shape that has come to connote search in both physical as well as digital spaces. Through referencing a magnifier, I hoped to convey a sense of the locator's functionality and to indicate how it should be oriented and held.

The locator's presumed RFID implementation dictated the need for a large flat surface area (to house the coil antenna), and the project's AR orientation demanded that this surface either be clear or hollow, so as not to obstruct vision. (If the locator obstructed vision, it would effectively replace rather than complement an existing way to search.) Since the power requirements of an RFID implementation would be comparable to those of a 1/10th scale radio controlled racing car (see calculations in Appendix F), the handle supporting the antenna was made large enough to accommodate a battery pack of the sort commonly used to power such vehicles. The battery would easily be the locator's heaviest single part, and placing the battery in the handle of the locator would help to ensure balanced handling. These technically dictated requirements – a large flat and clear surface area for an antenna and a handle proportioned for a long flat battery pack – reinforced the choice of a magnifying glass-like geometry.

In addition to communicating functionality and supporting an underlying implementation, the form of the simulated locator had to express provisionality. If this form were too polished, domain experts and other demo participants would be reticent to volunteer their own interpretations regarding what the hand-held locator might be *for*. Experimental participants would focus their attention on the issue of the locator's form, rather than attending to the more central issue of its behavior. To express provisionality, the locator's housing was constructed from cheap everyday materials using quick angular cuts. Cardboard laminates formed the base, while the antenna surface was constructed from clear PVC sheet stock and electrical tape.

In addition to choosing a form, I had to decide where to direct vibrotactile display on the user's device-holding hand. I decided to position the vibrotactile transducer under the thumb; since single thumbs oppose multiple fingers, a thumb is assured contact while in the act of grasping a handle. Directing vibration to the pad of one digit (rather than to multiple digits, or the hand as a whole) enabled controlled delivery of the vibrotactile stimulus, and made it possible to draw from an extensive body of psychophysical research concerning the sensitivity of the digit pads of human hands.

For demonstrations and experiments, the locator would require a button. This button was placed opposite the vibrotactile transducer, where it could be activated by the index finger like a trigger. (For a photograph depicting the position of the button and vibrotactile transducer, see Figure 22).

4.3.3. Lessons Learned

The process of developing and testing the simulation yielded numerous lessons. Some of these lessons influenced the design of subsequent demos and experiments, while others suggested alternate avenues of inquiry.

Informal testing of mapping functions (see Figure 23 for examples) suggested that some functions facilitated guiding the locator more than others. The difference between attempting to locate through discrete functions (e.g. Figure 23, upper left) and continuous, non-decreasing functions (e.g. Figure 23, upper right and bottom row) seemed more dramatic while the differences between locating through two alternate continuous non-decreasing functions were not so readily apparent. During informal testing, Gaussian and Sigmoid functions (Figure 23, lower left and center) were identified as two functions that seemed particularly well suited to assisting location. I hypothesized that this was because their slopes essentially conveyed three regions: “far”, “getting close” and “close”.

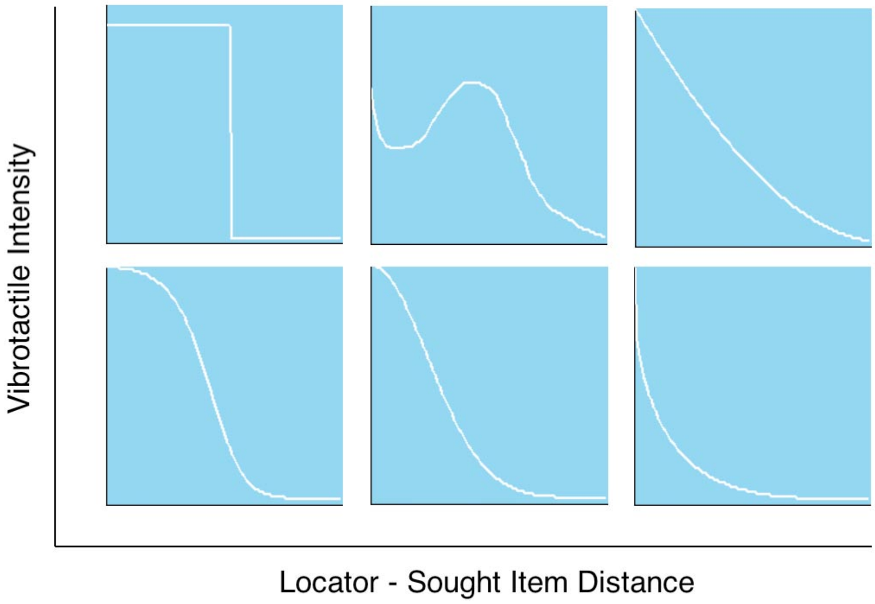

Figure 23: Some of the many functions used to informally explore mappings between locator-to-sought-item distance and vibrotactile amplitude.

Informal testing also revealed that thoroughly isolating vibration to the thumb pad would be a difficult technical challenge in its own right. A metal bar was initially placed in the handle to dampen vibration and model the weight of a battery pack, however this large mass of metal disrupted operation of the Fastrak sensor and so was removed. No non-metal weight was ultimately included.

Another potential problem identified during informal testing was the weight of the wire connecting the locator to the rest of the simulation’s apparatus. Informal testing revealed that the weight exerted an influence on how the locator could be held and moved. Since removing the wire was infeasible, participants were asked to “wear” the wire in a way that minimized its impact on arm movement. (For an illustration of this adopted posture, see Figure 24).

Figure 24: Posture participants assumed to compensate for the weight and stiffness of the locator's cord.

Early testing of the simulation also revealed that kinesthetic memory and spatial references would play greater roles in the act of locating than initially anticipated. Testing quickly revealed that searching a space through a vibrotactile gradient is not simply an activity conducted in the present, but an activity informed by past experience as well. Remembering where one has already searched informs where one decides to search next, and the spatial references provided by tables and shelves (as well as proprioception) serve to support this remembering. The value of spatial references was most evident when external spatial references were removed; the simulation permitted suspending sought item positions in “free space” far above tabletops, and such placements seemed significantly more difficult to find. This added difficulty was probably not solely due to the lack of visual references around the sought item; two other contributing factors were the absence of visual and passive haptic feedback from the (virtual) sought item itself, and the fact that “gravity-defying” objects are unusual – not a feature of everyday experience.

Since the sound channel of a PC drove the locator's vibrotactile transducer, informal testing of the simulation provided ample opportunity to compare the relative merits of audio and vibrotactile feedback. Since slight variations in amplitude were easier to detect through hearing than through touch, amplitude-modulated audio feedback clearly facilitated more sensitive guidance of the hand than amplitude-modulated vibrotactile feedback. Unfortunately, audio feedback also proved a disruptive annoyance for colleagues working in the same room (as well as adjoining rooms)[28]. Audio feedback (or combined audio/vibrotactile feedback) might ultimately prove to be more useful for steering a hand-held locator, however I chose to continue pursuing vibrotactile feedback for its non-disruptive nature.

4.4. Demonstration

4.4.1. Purpose

After developing the simulation and demonstrating it informally with colleagues and friends, I conducted a more polished demonstration. The aims of this demo were a) to receive feedback from a larger more critical audience, b) to verify that the simulation's apparatus was stable enough to support rigorous experimentation and c) to inform the design of experiments. It was important to obtain as much practical experience sharing the simulation as possible before designing experiments, so as to develop a sense of which aspects of vibrotactile targeting most warranted detailed examination. Laboratory experiments distill extremely complex situations down to much simpler ones in order to obtain tight experimental control, and so it was important to choose experimental variables with a global sense of what mattered most.

---

[28] Conceivably, audio feedback could be delivered privately through earphones, but earphones might destroy the ease with which the locator could be picked up and put down during the act of looking for a sought item.

4.4.2. Setting

Eurohaptics ’03, a conference hosted by Trinity College and Media Lab Europe in Dublin, provided the ideal setting for demonstration. This conference brought together hundreds of researchers from academia and industry to discuss theories of haptic perception and applications for active haptic display. One afternoon of the conference was devoted entirely to demonstrations, and during this afternoon session conference delegates were invited to wander the demo room floor at will, ask questions, and try out the devices on display.

The demo session provided an opportunity to introduce approximately fifty participants to the concept of locating through vibrotactile feedback, and to observe as they reacted to it in simulation. Due to the nature of the conference, most demo participants were familiar with active haptic display and used to considering hybrid physical/virtual constructions. They were predisposed towards understanding new applications for vibrotactile feedback, and professionally inclined towards offering insightful criticisms and suggestions. Unlike colleagues at the University of Limerick who had watched the project develop over a period of months, demo participants at Eurohaptics encountered the idea of locating items through a synthesized vibrotactile gradient all at once and for the first time.

4.4.3. Demo

The demo situated the simulation within the task of locating a book on a well-populated bookshelf, a task that the conference delegates might find somewhat familiar (see Figure 25). Locating an arbitrary book on a shelf was still, in a sense, a “toy” task – sought books are by definition not arbitrary – and this allowed participants to encounter the situation as the placeholder it was, rather than misinterpreting it as a tightly-defined and refined end application.

Figure 25: Demo participant getting a feel for locating books through vibrotactile feedback. (Photo taken after the demo).

Participants were told the aims of the project, showed how to hold and operate the locator, and given the opportunity to explore locating through a vibrotactile gradient for themselves. The order of these three activities was varied, so as to obtain a richer picture of how people made sense of what they were doing and feeling as they moved the locator and sensed the vibrotactile gradient. Simulation parameters were adjusted at various times throughout the demonstration (in no systematic way) to communicate to participants some of the dimensions of the design space under exploration, and to receive their feedback concerning what dimensions should be varied to best support the targeting of sought items.

4.4.4. Lessons Learned

The demo at Eurohaptics served to confirm design choices, suggest several different directions for experimentation and unearth problems for resolution.

The demo provided confirmation for many of the project's choices thus far. The metaphors used to explain the idea of a locator to colleagues at the University of Limerick (dowsing rod, Geiger counter, "hot/cold" children's game, etc.) also served well within the wider context of the Eurohaptics demo. The locator's physical form allowed demo participants to pick it up and wield it naturally with no previous experience, and did not draw attention to itself. The underlying implementation functioned without mishap during the two hours of the demo session. Participants often detected differences in the difficulty of locating books when simulation parameters (such as the mapping function) were altered, and expressed curiosity regarding the extent of these differences. Informal testing had suggested that spatial references within the search space (e.g. shelf lines) would influence searching behavior, and the demo confirmed this dependence; participants scanning motions were strongly defined by the shelves' lines and surfaces, in addition to the vibrotactile feedback. The demo confirmed various aspects of the simulation that had formerly rested on speculative, hypothetical and theoretical grounds.

In addition to confirming previous design decisions and observations, the demo suggested future directions for exploration. Some ideas cultivated from demonstration directly informed the design of experiments, while others suggested alternate paths of exploration.

One idea that emerged from the demo was to explore the relationship between the size of the sought item and the maximum range of the mapping function. Though a rough range for sought items had tacitly been assumed through choice of exemplars (e.g. books, hand tools, kitchen implements, key-chains, bottles, musical media, hats, etc.) the borders for this size had not been defined, nor its relationship to other parameters of the simulation considered. Through discussion with demo participants, it began to appear that the *ratio* of the sought item size to the range of the mapping function had an influence on the ease with which the sought item could be targeted. If the mapping range were too small and the sought item too large, the vibrotactile gradient would be of little assistance; finding the vibrotactile gradient would effectively mean finding the sought item. Conversely, if the mapping range was too big and the sought item too small, the gradient would be hard to detect. In this case, finding the target would be like trying to walk to the highest point on a salt flat. It seemed that the vibrotactile gradient would prove most useful in some "middle ground" between these extremes, where the ratio between the size of the sought item and the range of the mapping function was neither too large nor too small. (The extremes and middle ground of this ratio are illustrated in Figure 26).

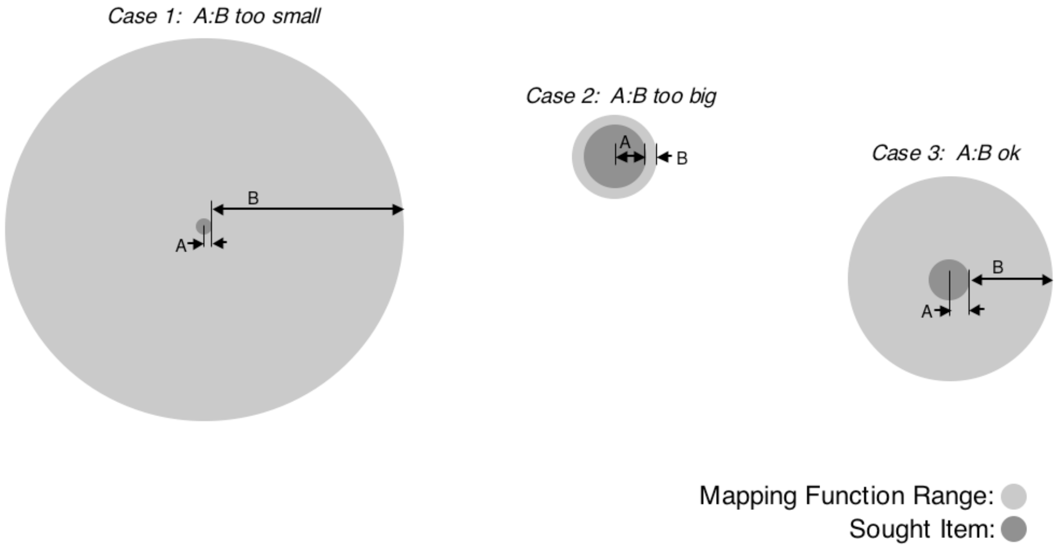

Figure 26: Ratio of Sought Item Size to Mapping Function Range. If the range of the mapping function is too large and the size of the sought item is too small (Case 1), the gradient is diluted by distance and can't be followed. If the range of the mapping function is too small and the size of the sought item is too large (Case 2), the locator is practically touching the sought item when vibrotactile assistance begins. Targeting assistance is likely to be most useful in the middle ground (Case 3), where the ratio between mapping function range and sought item size is neither too big nor too small.

Since the maximum range of the mapping function was a "soft" parameter (ultimately limited by the range of the presumed underlying RFID implementation), it could potentially be set *in reference to* the size of the sought item in order to achieve a desirable ratio. Sought items could potentially broadcast their size to the locator, and the locator could then use this size to determine a ratio between item size and mapping range that would facilitate targeting. (For this scheme to be feasible, an item would have to be associated with a sense-able measure of its own size – not a major technical hurdle given the presumed implementation of RFID and one additional assumption: that tags on items are programmed to report a measure of item size, perhaps as part of an automated manufacturing process.) Adjusting the ratio between item size and mapping function range seemed an intriguing, potentially practical and useful way to tailor vibrotactile location cues to sought items of different sizes[29].

Another idea resulting from the demo was the idea of accommodating more than one sought item. If the hand-held device could associate multiple items with multiple vibrotactile patterns, it could potentially serve not only as a locator but as a browser as well. Such a device might facilitate (for example) a librarians' search for out-of-place books on a shelf, or a pharmacist's search for all the bottles that had arrived from the west-coast supplier last Tuesday. While the possibility of supporting multiple sought objects seemed interesting, it also seemed far beyond what I could investigate thoroughly within my limited time frame.

A third idea that emerged from demonstration was to use the apparatus not as a crude simulation of an eventual RFID implementation, but as an interface (of the VR approach) in its own right. A few demo participants pointed out that the apparatus might be used as a means for placing and retrieving virtual objects, and suggested that this might prove useful in situations where visual display space was limited and virtual storage space around a user was desired. (For example, a hundred single-purpose virtual tools might be "hung" in space on virtual hooks, and retrieved through vibrotactile feedback when required). This idea suggested a very different trajectory for the project, and so it was not pursued.

In addition to suggesting directions for experimentation and confirming design decisions, the process of demoing unearthed two problems to be considered. One

---

[29] Gershenfeld, N. (1999) *When Things Start to Think,* Henry Holt & Company, New York. presents the notion that hybrid digital/physical objects should be designed as though they had "rights and responsibilities" in interaction. The expectation that sought items report their scale seems in line with this orientation.

problem related to the demo as a communication, the other related to the prospect of designing controlled experiments.

The demo did not effectively communicate the range at which location might realistically be assisted. When demo participants suggested applications, they tended to suggest applications that required a range far beyond that dictated by the constraints of the presumed underlying implementation. Over and over, it had to be explained to slightly crestfallen participants that the anticipated range for a functional system was on the order of a few meters, not even five to ten meters. Ideally, the demo itself should have made this clear without explanation, so as to elicit more viable applications.

The demonstration also revealed various physical and human factors that would make the design of experiments a challenge. Observation of the behavior of demo participants made it clear that the targeting of books through vibrotactile feedback was not only influenced by the spatial references provided by the bookshelves, but also by the books *on* the shelves. Participants paused to examine some books visually, but did not pause to examine others. Upon being found, sought books that were thick provided a greater margin of certainty than sought books that were thin. Sought books on the ends of shelves seemed easier to find than sought books towards the middle. In addition to uncovering physical variables pertaining to books and shelving, demoing revealed numerous human factors requiring consideration. Personal variation, interruption, motivation, fatigue and novelty would all have to be contended with in experimental design. These physical and human factors seemed extremely situated in nature, and drove home the realization that designing experiments that could at once claim tight experimental control, realistic modeling and relevance to real world situations – would be an extremely difficult challenge.

The demonstration at Eurohaptics yielded critical feedback from numerous domain experts, verified the robustness of the simulation, and provided a glimpse of numerous issues that would surface in the design of formal experiments. The positive reactions of demo participants provided additional confirmation that the idea of locating through vibrotactile feedback might be worth exploring.

4.5. Parameterization and Idealizations

After articulating the idea of vibrotactile location through scenarios, ensuring its feasibility through technical exploration, experiencing it personally through simulation and getting extensive feedback from others through demonstration, the investigation’s focus turned toward the design of experiments. Demonstration had supported the hunch that alterations to the distance/amplitude mapping could influence ease of targeting; experiments held the potential to clarify how, and how much. Before experiments could be designed, however, it was necessary to choose which idealizations could (and would have to) be made, and which parameters in the simulation most warranted in-depth examination.

4.5.1. Idealizations and Simplifying Assumptions

In order to achieve tight experimental control and ensure experimental feasibility, several idealizations and simplifying assumptions were necessary:

Targeting time ~ Targeting ease: It was assumed that the time spent trying to find a sought item was a reliable indication of the ease with which that item could be found. This assumption seemed reasonable, given the directed nature of searching to retrieve. The assumption appeared to hold true during demonstration, and went hand in hand with the early decision to direct experimentation towards a quantitative analysis at the level of human sensor-motor activity. (The assumption that time could serve reliably as a performance metric is intertwined with other assumptions typical of human factors experiments, i.e. that people understand a task they are asked to perform, and that motivation, distraction, fatigue and learning are either invariant between trials or can be compensated for. These concerns are discussed along with experimental design in the next chapter.)

In choosing to equate targeting time with targeting ease, one must recognize that time is not a foolproof metric, and that other metrics might be used instead of or in conjunction with time. Waving a locator around wildly within a search volume may find a target more quickly, but perhaps less easily. Additional metrics, such as the length of a path traveled by a locator on the way to a sought item, might be employed to for a more complete measure of "targeting ease".

Spherical, virtual sought items: Sought items were idealized as virtual spherical "targets". The choice to make sought items (targets) virtual enabled them to be placed automatically by software, and prevented tangible form from influencing hand motion during the act of locating. The choice to approximate targets as spheres made it possible for the Fastrak-based simulation to serve as a basis for an experimental apparatus. (Since the simulation depended on a particular Fastrak that could report position but not orientation, the simulation could support the placement and finding of shapes with radial symmetry more readily than shapes that lacked radial symmetry.)

Radial symmetry in vibrotactile feedback: The vibrotactile "active regions" surrounding sought items (targets) were, like the sought items themselves, idealized as spherical. This idealization unfortunately meant deviating from a realistic model of the presumed underlying implementation – field lines around an RFID tag are *not* radially symmetric – however the idealization was necessary due to the limitations of the simulation's underlying implementation (i.e. the Fastrak).

A specially-shaped "shelf space": An idealized, abstracted shelf-sized space was constructed instead of using an actual shelf for several reasons. First, the spatial references of shelves had influenced how people moved the locator in the demo, and I wanted to control for this influence in the experiments. Second, I was concerned that the distance between the locator's starting point (or "home" position) and the target positions might influence targeting time, and thought this could best be controlled through use of a custom search space. Third, I wanted to ensure that searching the space would require arm movement, but not walking, bending or twisting.

Suspend targets in space: Targets were suspended within the idealized search space (rather than placed on surfaces), because suspension prevented shelf geometry from becoming a confounding variable, permitted free movement of the locator within the space, and did not constrain the distance between the locator's home and target positions.

These idealizations preserved certain situation-independent aspects of locating through vibrotactile feedback (such as the rough scale of search space and sought items, the size of the locator, the manner of arm movement and the presence of a guiding gradient) and removed other, more situation-dependent aspects (such as the sought items' discrete physical presence, form, color and markings, and shelves' dimensions and placement).

Since locating under idealized conditions differs from locating under real conditions, it is useful to consider whether the over-all effect of the idealizations would be to make locating easier or more difficult. The situation-dependent aspects of locating that were removed through idealization would all likely make locating easier; thus the experimental bias was towards greater difficulty rather than greater ease. This bias seemed preferable, since it held the potential to accentuate rather than attenuate any observable trends of difference within experimental results.

4.5.2. Choice of Parameters

Since it was not feasible to examine the effects and interactions of all of the simulation's adjustable parameters (amplitude resolution, distance resolution, maximum range, target size, psychophysical exponent, mapping function etc.) in controlled experiments, a few parameters had to be selected as "most important" based on observations of the demo participants, personal experience and intuition[30]. The experimental variables ultimately chosen were: a) the mapping function and b) the ratio between target size and maximum range.

The choice of mapping function clearly influenced ease of targeting during the demos, and it seemed important to understand this impact. The mapping seemed a

---

[30] Conceivably, choice of experimental variables could have been informed by a preliminary sensitivity analysis of some sort. For lack of time, I did not pursue such an approach.

natural focus for a study in interaction design, since it coupled perceived cause with perceived effect through a dynamic representation.

Like the mapping function, the ratio between target size and maximum range influenced ease of targeting during the demo. Exploring the nature of this influence seemed worthwhile because it was an influence that could be compensated for by leveraging the identification aspect of the presumed underlying RFID implementation. The notion that sought items could automatically tailor the locator's vibrotactile feedback to facilitate locating seemed intriguing and potentially useful.

4.6. Tight Formulation of Research Questions

With preliminary explorations complete and experimental parameters and idealizations chosen, it became possible to formulate my research questions in a tighter, more answerable form. These questions became, in the order of their perceived importance:

A. How does the choice of the amplitude/distance mapping function affect targeting time? Given a set of mappings functions, does one mapping function result in shorter targeting times than another for a given participant? Does one mapping function consistently result in shorter targeting times than other mapping functions across a group of participants?

B. Given one chosen amplitude/distance mapping function: For a particular target size, does one value for the function's maximum mapping range consistently result in shorter targeting times than other values, across a group of participants? If so, what can be said across subjects about the relationship between target size and the range that results in shortest targeting times?

The ordering of the questions has a bearing on how they might be answered; I chose to explore the influence of the mapping function on targeting time first (because I thought this variable would have the greatest influence) and after choosing one mapping function as “best”, to explore the influence of the size/range ratio. With limited time and resources, it seemed most feasible to investigate the influences of the experimental variables one after the other.

Since this investigation speculatively presumes that a hand-held locator will prove useful, the questions were *not* framed so as to compare locating with and without assistance from a hand-held locator. Instead, the questions were framed to address how differences in vibrotactile feedback might make locating with a locator more or less useful (assuming the performance metrics of time and accuracy).

In addition to the two primary questions above, my chosen approach to evaluation suggested a third:

C. Does Fitts Law for targeting (an established predictive empirical relationship that relates targeting time to amplitude of an arm movement and size of a visible target) apply to haptic vibrotactile targeting?

Fitts’ law (see Chapter 2, Section 5) is a common referent for much of the work I had encountered on manual targeting, and so it made sense to relate it to my results[31].

---

[31] If Fitts’ law were to effectively describe vibrotactile targeting, this would bring the law’s application full circle; from the physical-world targeting tasks of the 1950s through the digital-world targeting tasks of the 1980s, to a digitally-mediated targeting task conducted in the physical-world.

5. Experimentation

5.1. Experiment Design

In order to address the three research questions formulated in the previous section (i.e. *A. How does the choice of amplitude/distance mapping function influence targeting time? B. How does the ratio between target size and mapping range influence targeting time? C. Does Fitts' law hold for blind vibrotactile targeting?*) two experiments were designed. This section discusses their design along several dimensions. First, the overall structure of the experiments and the experimental variables are introduced. Next, the experimental environment is discussed in terms of its geometry and apparatus. Following this, the format for time trials is presented. Finally, details specific to each of the two experiments are discussed.

5.1.1. Structure and Variables

I chose to conduct two controlled experiments; one to answer Question A, the other to answer Question B. Both were designed to yield data that would allow commenting on Question C. The parametric formulations for the two experiments were identical: an A*B*C*D*E factorial design, where the dependent variables were: mapping function, mapping-range, target-size (radius), home-target distance and number of repetitions[32], respectively. (For a diagram of experimental parameters and objects, see Figure 27). What differed between the two experiments was the number of values each experimental variable would take on (in each experiment, some variables would take on just one or two values).

---

[32] One thing to note is that repetitions are not strictly identical. The experimental variables – including home-target distance – remain constant across repetitions, however the position of the target does not. The positioning of the target is discussed in section 3.2: Experimental Environment & Apparatus.

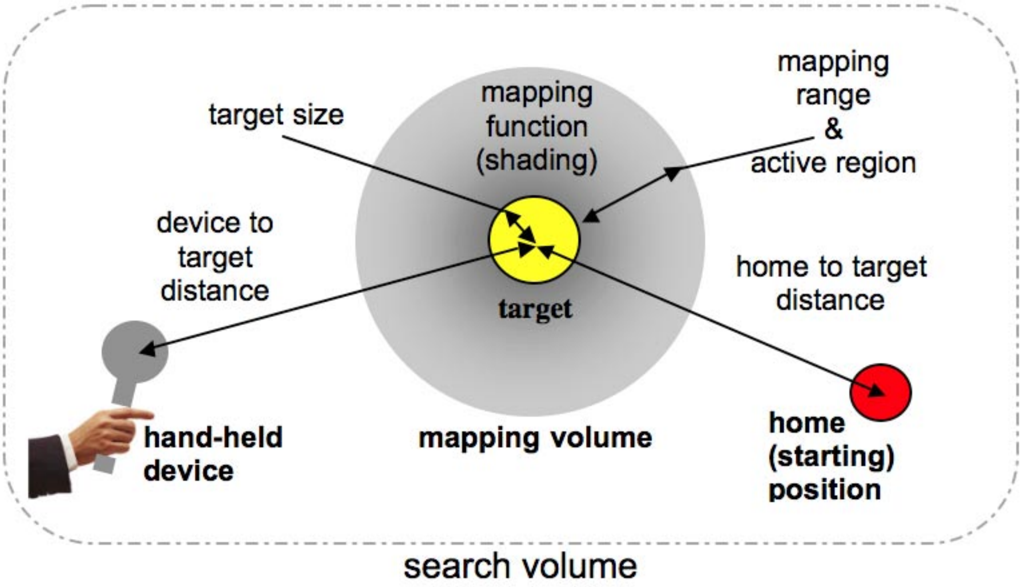

Figure 27: Diagram of experimental objects and parameters. The geometries of the search volume, home position and target are discussed elsewhere.

5.1.2. Environment and Apparatus – Design and Verification

The experimental environment (shown in Figure 28) consisted of a search volume, a “back end”, and a “rest station”. The search volume was designed to approximate a hemispherical region with a 0.5m radius, roughly equivalent to the length of a person’s arm[33]. The volume was raised off the floor so that its center was roughly level with a person’s chest. (Trials would be conducted with the participant standing). The volume’s concavity faced the participant, so that it would be easily accessible to the participant’s outstretched arm. The physical boundary of the search volume was visible, so participants could orient themselves with respect to it, and so that participants would not inadvertently search outside of it. The “home” position (starting position) for time trials was located at the center of the mouth of the hemisphere to allow for a variety of directions of movement during targeting. The home position was a point suspended in space, marked by a small but visible knot on an elastic piece of string stretched across the cavity of the search volume. This marker was necessary for participants to return to the home position accurately at the

---

[33] One could get carried away with justifying a particular arm length to use, and I didn’t want to become fixated on this biometric concern. I chose a radius of .5m based on feedback from people with different body proportions and the range constraint imposed by the presumed underlying implementation (RFID).

beginning of each time trial. Elastic string was used to hold the marker in place because it imposed minimal physical constraints on arm movement within the search volume.

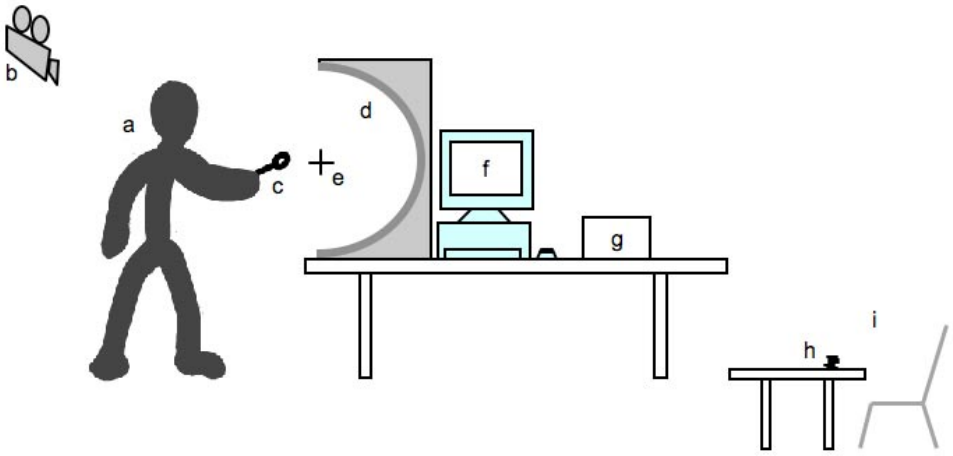

Figure 28: Experimental Environment & Apparatus: a) participant, b) video camera, c) hand-held device, d) search volume & support structure, e) home position, f) computer, g) Polhemus Fastrak 3D tracking system, h) Oreo cookies, i) rest station.

Target positions would appear (one per trial) as random distributions over hemispheres that were a) nested within the hemispherical search volume, and b) centered on the home position. The radii of the hemispheres upon which targets might appear were determined by the values of home-target distance. The rationale for introducing randomness in the placement of targets was to prevent participants from learning and remembering where targets were (or might be) over the course of the experiments. The rationales for fixing radii of target positions around the home position were a) to facilitate within-subject time trial comparisons, and b) to generate multiple time trial results for each home-target distance, results that might be used to explore the applicability of Fitts' law.

The experiment's "back end" was essentially the apparatus used for simulation and demos, with several software additions. One addition was the capability to automatically modify experimental parameters between trials. Another addition enabled the apparatus to keep track of time during the trials. A third addition made it possible to log experimental results – and the locator's path of motion – during each trial. Each of these additions to the simulation software reduced administrative overhead and made it possible to attend more closely to participants' actions while experiments were in progress.

The "rest station" consists of a chair and a plentiful supply of cookies. Participants were free to converse with others while resting.

Before using the apparatus for formal experiments, it was necessary to characterize it along dimensions of latency, linearity, and spatial resolution. If the vibrotactile output were not linear with the software-set amplitude of vibration, results would be confusing to interpret. If the latency of position sensing were too large, or its spatial resolution to crude, the apparatus would impede locating during the experiments. The test used to verify linearity appears in Appendix H. The system's input/output latency was determined, through the test described in Appendix I, to be 0.09 seconds – a delay comparable with the latency of the presumed underlying RFID implementation (0.1 seconds). Through comparing ruler-measured distances with Fastrak-reported distances, spatial resolution was determined to be within +/- 2mm within the search volume.

While conducting experiments in a controlled space apart from the distractions of everyday office life was desirable, it was not feasible. The experiments had to take place relatively near to where participants worked, so as not to inconvenience them. The experiments had to take place near an electronics lab, so that any technical difficulties might be addressed as they arose. The experimental environment had to be accessible at all points in the day over a period of several weeks. In light of these constraints, the best possible space proved to be a postgraduate student office, a space unfortunately subject to occasional distractions. The realized experimental apparatus and environment appear in the large photo of Figure 29 below. In this figure's main photo, the participant moves the hand-held locator in the search volume and the hidden target is visible as a circle on the monitor screen to the lower right.

Figure 29: Actual experimental apparatus and environment. Top: Sketches, paper models and construction of the search volume. (The idealized hemispherical search volume was realized as a multifaceted cardboard construction). Bottom: Participant within the experimental environment, engaged in locating a target in the search volume.

Participants were asked to wear earmuffs (rated for 27dB attenuation) during the experiment, so that the slight sound produced by the vibrotactile element would not influence targeting through the vibrotactile gradient. While combining sonic and haptic feedback would likely facilitate targeting (see discussion in Chapter 4, Section 3.3), these experiments would address targeting through vibrotactile feedback alone. (More accurately, active vibrotactile feedback plus passive visual feedback. Though completely blind haptic targeting would make for another interesting study, this was left for another time.)

I chose to remain present throughout the experiments in order to coordinate, observe, solve any problems that might arise, and discount trials where external interruptions or technical difficulties took place. Notes would be taken in writing on the experiments as they happened, and supplemented by video recorded from various angles.

5.1.3. Time Trial Format

The format for time trials was identical for both experiments, and consisted of the following steps:

1. The participant moves the locator to the home position (physically represented as a the knot on the elastic cord, digitally represented as a spherical region where vibrotactile feedback is turned off).
2. When the participant is ready to begin the trial, he/she moves the locator outside the home position's spherical boundary and the trial begins. A start time is recorded, and vibrotactile feedback is enabled.
3. When the locator enters the target region, the time of entry is recorded.
4. When the participant recognizes that the locator has entered the target (by sensing that the vibration delivered to his/her thumb has begun to pulse) and presses the button on the locator, a second time value – the time of perceived entry – is recorded. Also recorded are a) the position of the locator and target at the time of the button press, and b) whether or not the locator was in the target at the time the button was pressed. Pressing the button concludes the trial, regardless of whether or not the target is successfully found. (Note: participants could opt out of a trial at any time by pressing the button.)
5. The participant returns the locator to the home position when ready.
6. When the participant returns the locator to the home position, the values of experimental parameters are updated, and all is ready for the beginning of the next trial.

The time trials focused on "find-time": the interval between leaving the home position and pressing the button. Conceivably, the trials could have focused on other time intervals, such as the interval between leaving the home position and entering the target for the first time, or the interval between entering the target for the first time and pressing the button. I chose to concentrate on the interval between leaving the home position and pressing the button from the belief that entering the target (or its surrounding "active" vibrotactile region) multiple times might prove necessary; participants might build up a picture of the target's location iteratively, through employing memory and spatial abilities in combination. Multiple passes might reveal the target's location gradually, in the manner of an archaeologist taking a rubbing of a tombstone.

5.1.4. Experiment 1: Choice of Mapping Function

The first experiment was designed to answer Question A: to what extent does choice of the amplitude/distance mapping function impact targeting time? Naturally, the experimental variable of greatest interest for this experiment would be the choice of mapping function. Four different functions (illustrated in Figure 30) would be tested. The functions and their rationales were:

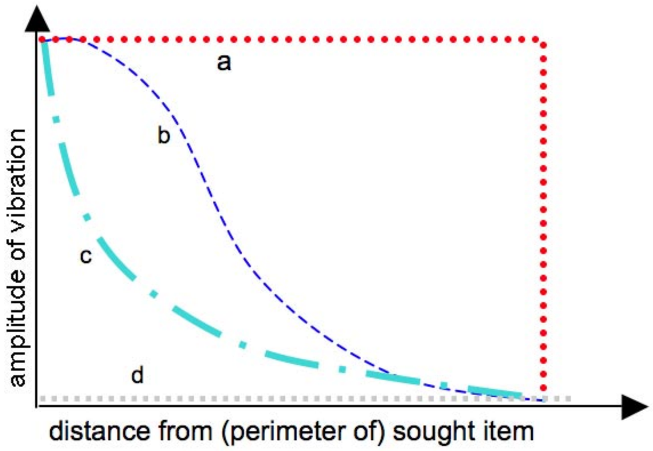

Figure 30: The mapping functions chosen for experimentation: a) Threshold, b) Perceived Gaussian, c) Natural d) None.

- None: The locator provides no feedback at all, until it is inside the target; this function was the control condition for the experiment.
- Threshold: The locator delivers feedback at maximum amplitude when the target is within the maximum mapping range, and delivers no feedback at all when the

target is outside the mapping range. This discrete, on/off mapping is, in a sense, the “baseline” for feedback provided by the vast majority of auto-ID (RFID & barcode) devices. Finding the target with the Threshold function, like the None function, entails a degree of random search.

- Natural: Amplitude of vibrotactile feedback is proportional to 1/(locator-to-target distance)^3. This mapping function describes the axial roll-off in signal strength of an RFID communication with distance. The function was included as one that would already exist naturally, underneath any artificial mapping function. If RFID signal strength were translated directly into the amplitude of a vibrotactile stimulus, this is the function that would best describe the relationship between distance and amplitude.
- Perceived Gaussian: Amplitude of vibrotactile feedback is proportional to a Gaussian-like function of locator-to-target distance (specifically, $Ae^{-cd\char`^2}$, where A is maximum amplitude, d is distance and c is a constant) all raised to the power of 1.72). The exponent of 1.72 compensates for the psychophysical complication that *perceived* vibrotactile amplitude is proportional to *actual* vibrotactile amplitude$^{(1/1.72)}$, and thus allows a person using the locator to *feel* a Gaussian roll-off. This compensation was included with the Perceived Gaussian function in order to foster continuity between what could be *felt* through vibrotactile feedback, and what could be *observed*/*changed* through the simulation’s GUI (i.e. the graph of the function). The value employed for the psychophysical exponent was that published by Franzen (1969) and Kenshalo (1978). Of the four functions, this was the one ostensibly designed for success; its gradient and magnitude both provided clues as to whether the locator was far away from, steadily approaching, or about to hit the target.

Ideally, other mappings (sigmoid, linear, etc.) would also have been included in this list, however the need to keep the duration of the experiments short precluded this.

While the mapping function was the variable of central interest in the first experiment, the range of the mapping function was also varied. The extent to which one mapping function might facilitate speedier targeting than another other might depend on the range of the mapping function, and to reveal this possible dependence, two values for the mapping-range variable were employed. The specific values, 18cm and 32cm, were selected through informal testing to fit comfortably within the ratio limits illustrated by Figure 26, and the range limitations dictated by present-day RFID regulations and technology.

Only one value, 6 cm, was used for the target-size (radius) variable. A value much smaller than the overall dimensions of the search volume was chosen to ensure that finding the target would be a non-trivial task.

Home-target distance took on two values: 20 and 32cm ("short" and "far" within the search space, as determined through informal testing). The number of repetitions was set at 6, so that trials could be discounted if necessary and averages might be taken. In summary, the first experiment's factorization became:

[mapping choice]*[range]*[target-size]*[home-target distance]*[repetitions]
4 * 2 * 1 * 2 * 6

This brought the total number of trials to 96. At 10 seconds per trial, the total time required for one participant to complete the experiment was estimated to be approximately 16 minutes.

Participants would be selected from the available pool of graduate students: men and women between the ages of 20 and 40 who reported being comfortable with new computer-related technologies. Early adopters were desired for their enthusiasm, patience with technical difficulties, and readiness to articulate possible end applications for a hand-held locator.

Before the experiment, each participant would be briefed on the concept of a gradient, the sensation of vibrotactile feedback, and the idea of locating through a gradient in vibrotactile feedback. Participants would be allowed 2 minutes of practice time – just enough to get the idea and get "warmed up" for the trials.

After this practice, the experiment would begin. There would be two intermissions during the experiment to reduce the effects of fatigue and allow opportunities for questions and comments. The intermissions would last approximately 2 minutes and will occur 1/3 and 2/3 of the way through the total number of trials. (For the script I would use during the briefing and experiment, see Appendix J). When the participant completed the experiment, he/she would be asked to fill out a questionnaire (see Appendix K). For the Declaration of Informed Consent test participants would be asked to sign, see Appendix O.

The number of participants was chosen to be between 25 and 50; tests would continue until at least 20 subjects had agreed to participate in the second experiment. (The questionnaire would request this participation).

My predictions regarding the outcomes of the first experiment were as follows:

- The choice of mapping function will make a difference in targeting time.
- The Threshold mapping function will not allow quickest over-all targeting.
- The Perceived Gaussian mapping function will allow quickest over-all targeting.
- The benefit of the “best” mapping function over the others will be less pronounced for the smaller of the two values for mapping range.
- The rank of “best” to “worst” mapping function will be independent of the value of mapping range.

5.1.5. Experiment 2: Target Size and Mapping Range

The purpose of the second experiment was to answer Question B: How does the ratio of target size to mapping range influence targeting times, given the mapping function selected as “best” after Experiment 1? The experimental variables of greatest interest in this experiment were target-size and mapping-range. (For a diagram illustrating these parameters, see Figure 27).

A number of factors constrained the selection of values for target-size and mapping-range: the need to keep the experiment short, the size of the search space and the desire to generate results that complemented the results gathered in Experiment 1.

Since the size of the search space would remain constant, increasing the radius of the target would increase the proportion of the search space that *was* the target. To minimize the effect this would have on targeting time, the values used for target radius, 4cm and 6cm were chosen (through informal testing) to make the target much smaller than the dimensions of the search space.

The need to keep the experiment short constrained the number of values that target size and mapping range could take on. Since the target had to be kept small, target size would only assume 2 values, while mapping range would take on 4. The values were chosen through informal testing so as to fit comfortably within the ratio limits illustrated by Figure 26. One of the two target-size values and two of the four mapping-range values were copied from the first experiment, so that trial data from the two experiments might be comparable for a given participant. The 2 remaining maximum mapping range values would be chosen after the first experiment had been completed. (The first experiment will clarify whether it makes sense to focus on larger or smaller values of mapping range during the second experiment). The remaining value for target size would be chosen so as not to alias target size : mapping range ratios.

Two values would be used for home-target distance. They differed from the two home-target distances selected for use in the first experiment, so that by the end of Experiment 2, comparable data sets would exist for 4 distances. This staggering was chosen to facilitate answering Question C (i.e. Does Fitts' law apply?).

As in Experiment 1, the number of repetitions was set to 6. In summary, the second experiment's factorization would be:

$$[\text{mapping choice}]*[\text{range}]*[\text{target-size}]*[\text{home-target distance}]*[\text{repetitions}]$$
$$1 * 4 * 2 * 2 * 6$$

This brought the total number of trials to 96, as in Experiment 1. At 10 seconds per trial, the total time for the trials was estimated at approximately 16 minutes.
The participants for this experiment (>20) would be selected from the pool of participants who participated in the first experiment. The total number of participants would be between 20 and 30. As in the first experiment, there would be two intermissions of approximately 2 minutes that will occur 1/3 and 2/3 of the way through the total number of trials. Upon completing the experiment, participants would be asked to respond to a questionnaire. (For the script I would use during the experiment, see Appendix L. For the questionnaire, see Appendix M.)

My predictions regarding the outcomes of the second experiment were as follows:

- The ratio of target size to mapping range will have an effect on targeting times.
- Too large a ratio of target size to mapping range will be more of a problem (will result in longer targeting times) than too small a ratio of target size to mapping range.
- Fitts' law will not hold for vibrotactile targeting.

Table 1 summarizes the variation of variables within the first and second experiments.

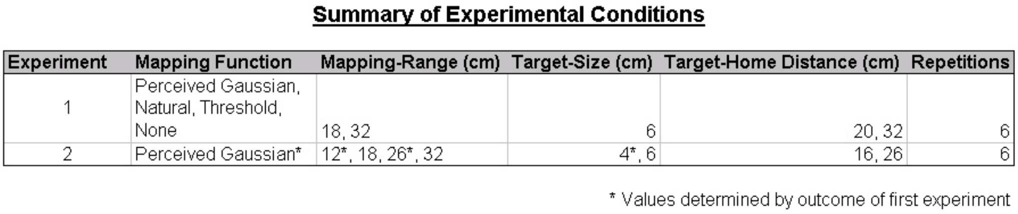

Summary of Experimental Conditions

Experiment	Mapping Function	Mapping-Range (cm)	Target-Size (cm)	Target-Home Distance (cm)	Repetitions
1	Perceived Gaussian, Natural, Threshold, None	18, 32	6	20, 32	6
2	Perceived Gaussian*	12*, 18, 26*, 32	4*, 6	16, 26	6

* Values determined by outcome of first experiment

Table 1: Summary of the experimental variables and their values for each of the two experiments.

5.2. The Experiments

The experiments were conducted over a period of five weeks; each experiment took two weeks, and there was a one-week pause in between. This intermediary interval was relatively short because the Fastrak position-sensing system was required by other projects, however it proved sufficient for selecting values for the parameters of the second experiment that depended upon completion of the first (the mapping function, and one of the two values for target size).

Relatively few “in-flight” corrections were necessary during the experiments. These concerned preparing the participants and minimizing distraction within the test environment.

Early into conducting the first experiment, it became clear that one of the most error-prone parts of the trial cycle was the moment at which participants pressed the button to end the trial. At this moment, participants displayed a propensity to move the locator slightly while pressing the button. Often, this movement had the unintended effect of pushing the locator out of the target just when it should have been recorded as being inside the target. (This movement could be noticed by keeping one eye on the on-screen representation of the target and locator, and the other eye on the experimental participant.) The glitch occurred so often that I chose to post-process participants’ result files based upon the following rule: If the trial was reported as a miss but the locator had been in the target during a 2-second window just prior to the button-pressed, the miss was re-recorded as a hit. This post-processing may inadvertently have changed a small fraction of legitimate misses to hits, however the benefit of changing all the should-have-been-hits into hits was worth this cost.

Anther error-prone moment in the trial cycle was that of returning the locator to the home position. Since this position was indicated by a “soft” physical constraint (a knot in an elastic cord stretched across the cavity of the search volume) it was possible for participants to return the locator to the home position, and then accidentally start the next trial by unintentionally moving the locator out of the spherical region that delimited the home position. Effectively, the passive feedback that participants received to signal “in home position” and “out of home position” was weaker than anticipated. This source of error was addressed through instruction, and once participants were shown how to move the locator into and out from the home position precisely, false starts were not a significant issue.

Initially, participants seemed apprehensive about touching or moving the elastic string while targeting. Pulling the string back and demonstrating its elasticity during the experiments’ initial training period served to alleviate participants’ apprehensions and encourage them to move naturally throughout the search volume.

Another tendency that participants displayed from the outset was a tendency to get “stuck” searching along the surfaces of the search volume’s boundary. This was anticipated, since the boundary provided spatial references for searching motions, and unavoidable (without the boundary, participants might search outside the volume, where the target was guaranteed *not* to be). To counterbalance this tendency to “cling to the walls”, I began to emphasize to participants during the initial training period that the volume had *depth*, that the target could be near the center of the volume, or near the periphery.

Some participants seemed to lose sight of – or else momentarily not care about – the fact that they were engaged in timed trials. Occasionally, participants would react to distractions, take impromptu breaks, or move the locator in exploratory (as opposed to goal directed) patterns, such as finding the target region multiple times from alternate directions, attempting to encircle the target region without entering it, etc. On such occasions, I reminded participants that the validity of the experimental results depended upon effort and attention, and asked them to report any trials where breaks proved necessary or distractions occurred. In addition to verifying that participants understood the expectations associated with time trial experiments, it was also necessary to ensure that colleagues in the area surrounding the experimental environments understood them as well. When experiments began, colleagues would occasionally interrupt participants in the middle of testing without realizing what the effects of interruption would be on the experiment. A posted explanation and an appeal for support garnered the necessary understanding. (Controlled experiments did run against the grain of the office space, so it was fortunate that colleagues were patient with the process).

Aside from these minor adjustments, I stuck to the scripts (Appendices J&L), observed the trials, took notes and shot video. Figure 31 displays stills of several participants engaged in targeting.

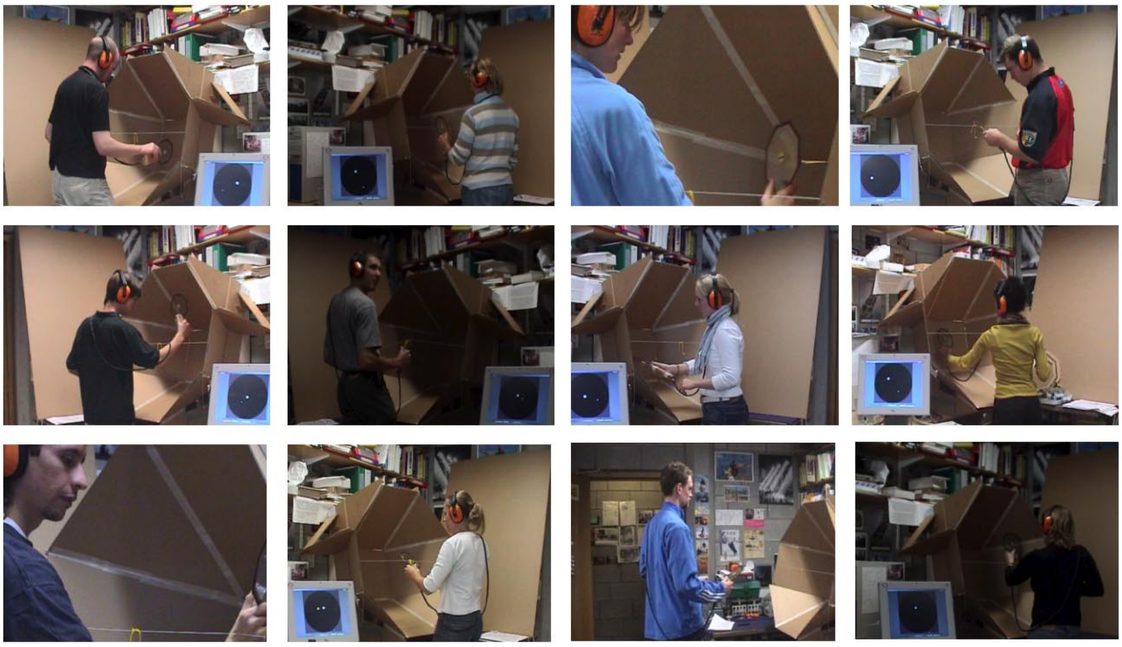

Figure 31: Participants in the midst of time trial experiments.

5.3. Results & Discussion

This chapter begins with a discussion of qualitative results, so that they may provide a commonsense "reality-check" on all the experiments' findings. Next, quantitative results are introduced, and discussed in relation to the three research questions articulated in Chapter 4, Section 6. Finally the impacts of learning, fatigue and other sources of error are addressed. The intent of this discussion is not to provide an exhaustive analysis – the data can certainly be viewed in other ways – it is to provide one coherent and rational traversal through an unfamiliar data space.

5.3.1. Qualitative Results

In order that the scope and limitations of the quantitative results may be fully understood, it is useful to first consider the qualitative results gleaned from the experimental process. The coming sections will first present observations at a general level, and then address the specific topics of the questionnaire.

5.3.1.1. General Observations

Participants encountered and made sense of the time trial experiments in various ways. For some, the experiments were a competition to be won. For others, the experiments were opportunities for contemplating an unusual phenomenon. Some participants clearly regarded the time trials as a game to be played, while others treated it as a chore to be completed.

The experimental apparatus and environment served well, though not without occasional hiccups. The simulation software and Fastrak stopped working occasionally, and some trial data was lost as a result. The test environment proved suitable most of the time, however there were occasional interruptions due to the public nature of the office space. The cardboard-delineated search volume seemed to accommodate participants of various body proportions well. The locator's form proved sufficiently comfortable for most participants, though a few participants with smaller hands complained that its grip was uncomfortably large for them.

The time that participants actually took to complete the experiments was much greater than anticipated; while the estimated duration for each experiment was on the order of 10-20 minutes, the actual time taken was closer to 50-60 minutes. This disparity was probably due to the fact that initial estimates had been based upon trial runs conducted by an experienced participant using the vibrotactile mouse and targeting in a two-dimensional search space, rather than a novice targeting with the Fastrak-based locator in the three-dimensional search volume.

5.3.1.2. Questionnaire Responses

This section summarizes participants' responses to select questionnaire questions, and supplements the responses with my own observations. Responses to the more demographic questions – regarding age, handedness, height, etc. – are presented quantitatively, in Chapter 5, Section 3.2.1. Additional response summaries are included in Appendix N.

First Questionnaire: Influence of the Mapping Function

How did you find the target? What strategies did you try?

Participants reported using a variety of techniques and strategies for targeting. Some participants steered towards the target slowly and carefully along linear trajectories. Others reported moving the locator wildly throughout the search volume at first, then successively refined their motions based on the vibrotactile feedback they received. Some participants reported initially feeling timid about making gross arm movements with large correction factors, but added that their inhibitions evaporated as the experiment progressed. Some participants reported experimenting with creative visualizations such a closing their eyes, and pretending they were looking for the nucleus in an atom[34].

Participants' experimentation with different strategies was manifest visually in different patterns of arm motion. Sometimes participants "swept" through the search volume from side to side, then from top to bottom in a long "rasterizing" gestures. At other times, participants moved from the center of the volume outward in concentric spirals, in the manner of a fencer parrying an attacker's foil. Another pattern of arm motion that participants displayed suggested a divide and conquer strategy: search near facets of the search volume boundary, then search the volume's central hollow. Participants tended to try several patterns initially, then settle on a favorite pattern for the remainder of the experiment with minor variations. Since the

---

[34] In reviewing questionnaire responses, it was interesting to note how people struggled to describe what they were experiencing; people wrote of "thumping dots". "atoms" "auras" and "parabolic buzz". Such descriptions would doubtless be useful to consider in future efforts to communicate the idea of targeting through vibrotactile feedback.

real-time position of the locator was recorded during the experiments, it proved possible to reconstruct and display patterns of arm movement. Some examples appear in Figure 32 below:

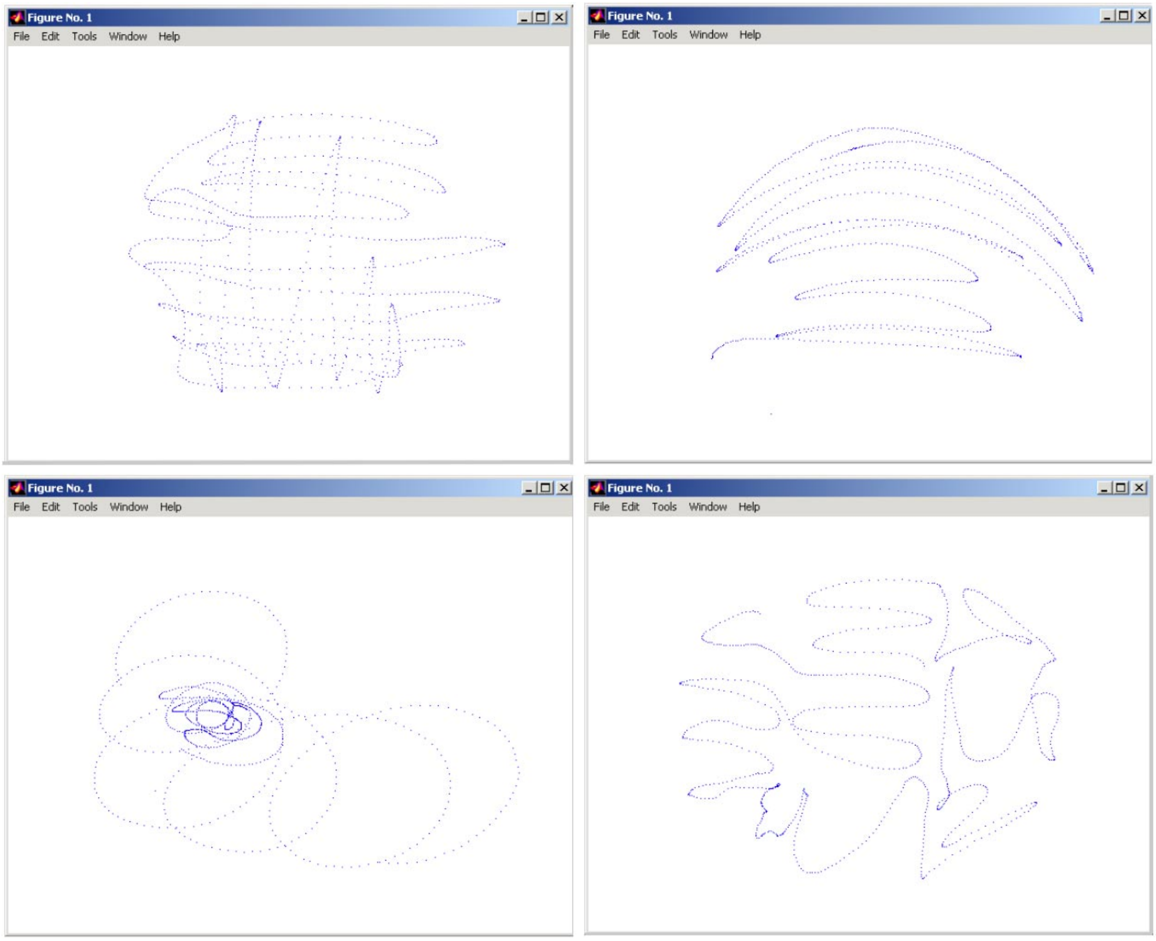

Figure 32: Reconstructed movement patterns of the hand-held locator, as wielded by several participants during the time trial experiments. (The plane of the pictures is parallel with the open mouth of the search volume). Tightly packed data points indicate slower motion, while loosely packed data points indicate faster motion. These plots illustrate something of the variety of techniques participants explored while targeting.

Some participants reportedly employed a “branched” strategy; they chose how they would move the locator based on the mapping function they were presented with in a given trial. The fact that participants relied on different strategies – and sometimes changed strategies with the experimental parameters – probably had some impact on the experimental results. Though the experiments had been conceived as a way to investigate low-level human sensor/motor activity, high-level psychological factors clearly came into play. A great deal of time and effort could have gone into identifying & categorizing strategies, and examining how and when people shifted between them; I chose simply to 1) note that there were different strategies 2) provisionally assume that their influence on trial times was much smaller than the influence of the experimental variables under scrutiny 3) and press on.

How many [patterns of vibrotactile feedback] do you remember? How were they alike/different?

Most participants recalled and identified three different mappings: one where there was no vibrotactile feedback (until the target had been found), one where feedback was constant within a certain range, and one where feedback varied continuously with distance. A few participants reported more than one continuous amplitude/distance mapping (there were 2). The participants' responses served as a reality check for the experiment; they suggested that a) there *were* detectable differences between the mapping functions, and that b) some differences between mapping functions were subtler than others. (It was assumed that participants' responses to this question were not limited by recall; the actual number of mapping functions – four – was relatively small, and each mapping function was delivered 24 times during the experiment).

Did you prefer one pattern of vibrotactile feedback to the others?

Participants reported an almost universal preference for the gradually graded mapping functions over Threshold or None functions. (The only participant who did not report this preference interpreted the question as referring to the difference between the 250Hz vibration and the 10Hz on/off modulation of the 250Hz vibration.) Most participants were aware of the experiment's aims and hypotheses, and may have reported a preference of gradual gradients in part to please the experimenter; the presence/extent of this possible bias remains unknown.

While the reported preference for gradual gradients was to be expected, participants' preference of no feedback at all to the Threshold mapping function came as a surprise. Perhaps, as one participant mentioned during the experiment, the Threshold mapping function was interpreted as being "taunting", in the sense that the system "knew" where the target was, and wasn't being helpful about it. Another justification that participants volunteered for their preference of no feedback at all over the Threshold mapping function was that the constant, full-on vibration of the Threshold function was uncomfortable, tiring and numbing. (Kenshalo, 1978) reports that subjective sensitivity to a vibrotactile stimulus can fall by 25% after 4 seconds due to habituation; the duration of many time trials in the experiment was much longer than 4 seconds.

Was this experience of trying to find a location through vibrotactile feedback like anything you've done before?

Several participants (7/25) reported that the experience was not like anything they had done before. The remaining participants associated targeting with using a metal detector, finding a light switch by feel in the dark, finding wall studs by tapping, stabilizing a kayak by feel, playing blind man's bluff, playing hunt the thimble and swinging a ping pong paddle or tennis racquet. (These associations might loosely be categorized in terms of finding the imperceptible, playing a game and making back-and-forth hand movements across one's field of view.)

In what situations or occupations could you envision this capability being useful?

Participants reported numerous situations where a hand-held locator might prove useful. These situations were grouped into the following *types* of situations:

1) Finding items within a certain size range for subsequent retrieval where:
    a) Vision is not available (e.g. in darkness, blindness)
    b) Vision doesn't help, because:
        - Containers occlude contents (e.g. boxes in attics, warehouses and mail rooms)
        - Items in a collection are visually homogeneous (as in the large collection of mix CDs depicted in Figure 33).

- The search space is visually busy (e.g. a messy room[35], an unfamiliar office shelf).
- The search criterion is not essentially a visible quality (e.g. sell-by dates in a supermarket, last check-out dates in a library, dates of arrival in a pharmacy)

Figure 33: Writable "mix" CDs – an example of homogeneous packaging that does not assist visual search.

2) Finding things you know you will misplace (e.g. keys, wallets, passports, hidden funds)[36].
3) Localizing imperceptible but dangerous items to prevent access/contact (e.g. chemicals in a laboratory, medicines that should be kept out of children's reach).
4) Games about unmasking spatial relations (i.e. treasure hunts, connect-the-dots in 3D).

Categorizing the types of situations that participants volunteered helps to define the borders of the various "problem spaces" one might search while working backward from the hand-held locator as a solution, towards an actual real-world problem. In general, participants did not volunteer applications that specifically required or leveraged *vibrotactile* feedback. In responding to the questionnaire question, they

---

[35] While the occupants of "messy" rooms often and amusingly know just where everything is, others may not and may need to know in certain situations.
[36] One participant noted that forgetfulness and misplacement are often issues faced by the elderly, and entrepreneurially hinted that the average age in countries such as the United States is on the rise.

focused on locating, rather than communicating location through any specific perceptual channel.

Second Questionnaire: Influence of Target-Radius-to-Mapping-Range Ratio

Did the range of the feedback change during the experiment?
If so:

> *How many different feedback ranges do you recall being presented with?*
> *Did you find it easier to locate the target when the range was larger or when it was smaller?*

The majority of participants (13/18) reported experiencing 3 or 4 different mapping-ranges. (There were 4). Most participants (10/16) reported that larger ranges made targets easier to locate, however a significant portion (4/16) reported the opposite. This ambivalence may suggest that the ratio of target-size to mapping-range (rather than simply the mapping-range on its own) does in fact exert an influence. On the other hand, it may reflect individual differences, the use of different targeting strategies, or the lack of a clear relationship between ease of targeting and mapping range.

Did the target change size?
If so:

> *How many differently sized targets do you recall being presented with?*
> *Did you find it easier to locate the target when the range was larger or when it was smaller?*

Most participants (14/16) accurately reported that the target did change size; roughly half of the participants reported two sizes, while the other half reported 3 (there were 2). Nearly all of the participants (14/16) reported that targeting was easier when the target was larger.

Based on your experience of this study, how would you recommend that this investigation proceed? (What would you see as the "next step"?)

After encountering the concept of locating through vibrotactile gradients in the experiments, participants were reportedly eager to try out the concept in a more concrete fashion; with actual physical items, a wireless locator and a real-world task. Additionally, participants were curious about the merits of other forms of feedback (e.g. audio and visual), and how these might combine (or replace) vibrotactile feedback to improve targeting. Like the demo participants at Eurohaptics, participants in the experiments wanted to see a system with greater range than the RFID-limited range that was emulated by the experimental apparatus.

5.3.2. Quantitative Results

5.3.2.1. Participant Demographics

There were 24 participants in the first experiment: 11 women and 13 men, with an average age of 29. The majority of the participants were in their twenties and thirties; the oldest participant was 52, while the youngest was 16. Participants' height varied from 191 cm to 155 cm, with an average height of 150 cm. Of the 24 participants, 21 were right-handed, 2 were ambidextrous and only 1 was left-handed. Age, height, gender and handedness data for each participant for both experiments appears in Table 2.

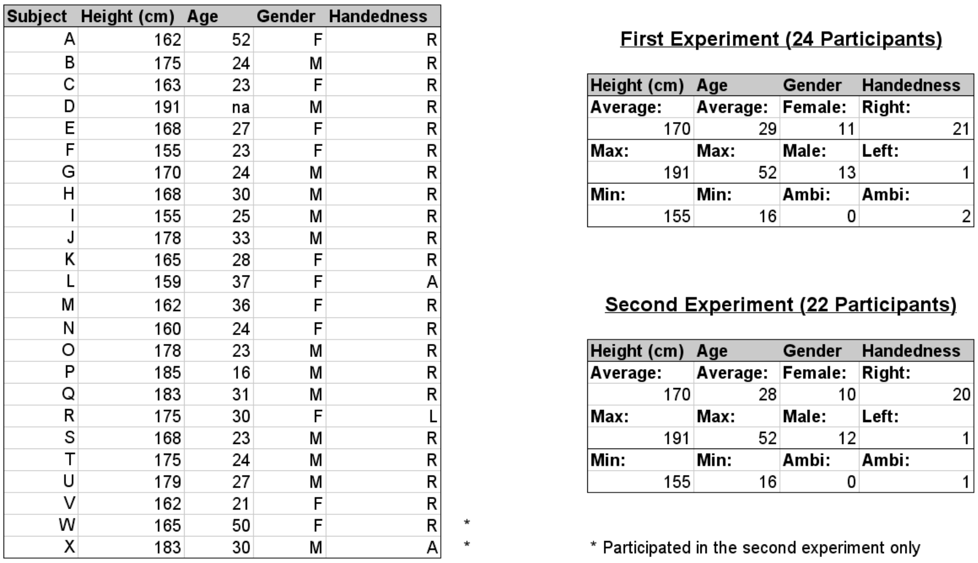

Subject	Height (cm)	Age	Gender	Handedness	
A	162	52	F	R	
B	175	24	M	R	
C	163	23	F	R	
D	191	na	M	R	
E	168	27	F	R	
F	155	23	F	R	
G	170	24	M	R	
H	168	30	M	R	
I	155	25	M	R	
J	178	33	M	R	
K	165	28	F	R	
L	159	37	F	A	
M	162	36	F	R	
N	160	24	F	R	
O	178	23	M	R	
P	185	16	M	R	
Q	183	31	M	R	
R	175	30	F	L	
S	168	23	M	R	
T	175	24	M	R	
U	179	27	M	R	
V	162	21	F	R	
W	165	50	F	R	*
X	183	30	M	A	*

First Experiment (24 Participants)

Height (cm)	Age	Gender	Handedness
Average:	**Average:**	**Female:**	**Right:**
170	29	11	21
Max:	**Max:**	**Male:**	**Left:**
191	52	13	1
Min:	**Min:**	**Ambi:**	**Ambi:**
155	16	0	2

Second Experiment (22 Participants)

Height (cm)	Age	Gender	Handedness
Average:	**Average:**	**Female:**	**Right:**
170	28	10	20
Max:	**Max:**	**Male:**	**Left:**
191	52	12	1
Min:	**Min:**	**Ambi:**	**Ambi:**
155	16	0	1

* Participated in the second experiment only

Table 2: Participants for the first and second experiments. Left: The full demographic data. Right: The summaries, by experiment.

Twenty-two of the 24 participants went on to complete the second experiment, so the participant demographics changed little between the experiments. There were 10 women and 12 men. The average age was 28, the average height 148 cm. Twenty of the 22 participants in the second experiment were right-handed, 1 was ambidextrous and 1 was left-handed.

All participants were familiar with desktop computers, and all had in common an enthusiasm for interactive systems.

5.3.2.2. Missed and Canceled Trials

The first step in analyzing time trial data was to examine the miss/cancel trials, the trials participants ended by pressing the button at a moment when the locator was not in the target. When the breakdown of "hits" and "misses/cancels" was tabulated and graphed for all of the trial data[37] from the first experiment, (~90 trials/participant for 24 participants; see the left-hand graph of Figure 34), I found that misses/cancels made up only 5% of the total number of trials. This suggested that misses/cancels would not have a strong limiting influence on sample size (an important consideration for statistical analysis), and that emotional/psychological issues such as discouragement and impatience would not have a strong bearing upon the experimental results.

When the total number of misses/cancels was separated by mapping function, the majority of misses/cancels corresponded to the Threshold and None mapping functions, while relatively few of the misses/cancels corresponded to the Perceived Gaussian and Natural functions. In decreasing order, the relative frequencies of misses/cancels for the four mapping functions were: None:60%, Threshold:26%, Perceived Gaussian:10% and Natural:4% (see right-hand graph of Figure 34). It was not surprising that the None function accounted for most of the misses/cancels. This was the experiment's control condition, in which participants were hunting for the target without assistance. That the Natural and Gaussian functions together accounted for a minority (14%) of the misses was also not surprising, since these functions provided more assistance than the Threshold function.

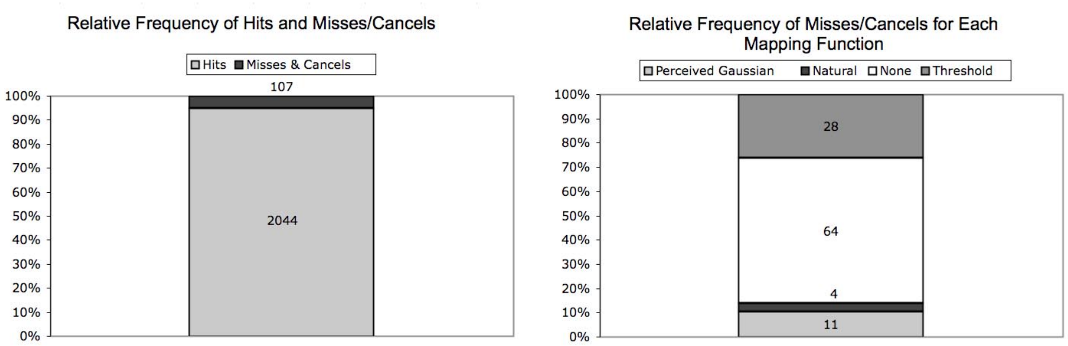

Figure 34: Left: Percentages of targeting hits and misses/cancels for all trials (all participants and all mapping functions included). Right: Relative frequency of misses/cancels for each mapping function (all participants included). (The number on or above each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

After the miss/cancel trials had been examined they were removed them from the data sets for both experiments. Subsequent analyses would focus upon trials in which the target was found.

5.3.2.3. Find-Time Distributions: Mean, Standard Deviation and Skew

Once miss/cancel trials had been examined and removed, the distributions of the find-time values for each participant's data for the first experiment were considered. How these trial time values were distributed would dictate which statistical tools would be appropriate to apply. In examination of the find-time data, three statistical parameters were considered: mean, standard deviation and distribution skew.

Means

For each participant, the mean find-time for all the (non-discounted, non-miss/cancel) trials of the first experiment were determined, then compared across participants. The means ranged from 12 to 33 seconds, with an average of 22 seconds. While the absolute value of participants' find-times may be meaningless

---

[37] All the trial data that was not discounted due to technical difficulties or obvious distractions

without an actual real-world task defined by participants themselves, the relative values were examined to see if they might unmask significant demographic correlations. Participants were ranked based on their overall mean find-time (see Table 3), and no clear differences based on age, gender, height or handedness were found. The three fastest participants were male and the three slowest were female, however the *next* three fastest participants were female, and the *next* three slowest participants were male. The overall mean find-times of the two ambidextrous participants mirrored across the median of the mean find-times, and the one left-handed participant's mean find-time was near this median of means. Height and age seemed evenly distributed across the scale from fastest to slowest, with the possible exception that the two oldest participants both appeared in the fastest quartile. One observed difference (that did not appear in this ranking) was that older participants seemed more willing to cancel trials; perhaps this reflected an attitude of greater detachment towards the targeting task.

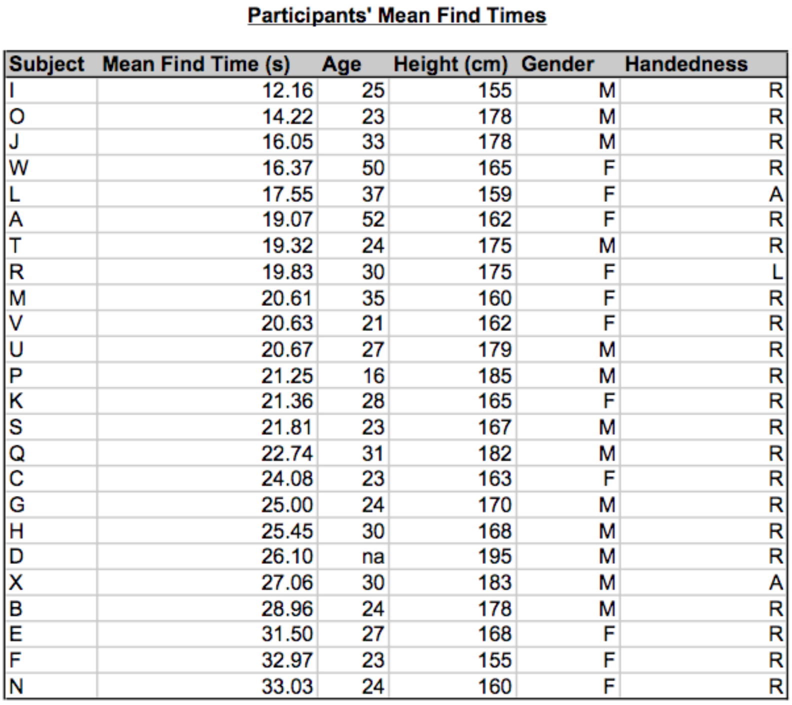

Participants' Mean Find Times

Subject	Mean Find Time (s)	Age	Height (cm)	Gender	Handedness
I	12.16	25	155	M	R
O	14.22	23	178	M	R
J	16.05	33	178	M	R
W	16.37	50	165	F	R
L	17.55	37	159	F	A
A	19.07	52	162	F	R
T	19.32	24	175	M	R
R	19.83	30	175	F	L
M	20.61	35	160	F	R
V	20.63	21	162	F	R
U	20.67	27	179	M	R
P	21.25	16	185	M	R
K	21.36	28	165	F	R
S	21.81	23	167	M	R
Q	22.74	31	182	M	R
C	24.08	23	163	F	R
G	25.00	24	170	M	R
H	25.45	30	168	M	R
D	26.10	na	195	M	R
X	27.06	30	183	M	A
B	28.96	24	178	M	R
E	31.50	27	168	F	R
F	32.97	23	155	F	R
N	33.03	24	160	F	R

Table 3: Participants, ranked by over-all mean find time for all trials from the 1st experiment. No clear correlations between mean find-time and age, gender, height or handedness are apparent.

Standard Deviations

After participants' mean find-times over all trials had been examined to obtain a sense of demographic differences, the ratio between mean and standard deviation for find-time distributions corresponding to each mapping function were examined. Dividing each participant's mean find-time for each mapping function into the mean's associated standard deviation, and then averaging the ratios across participants yielded average ratios of 0.97, 0.88, 0.99 and 0.93 for the Perceived Gaussian, Natural, None and Threshold mapping functions, respectively; the magnitude of mean and standard deviation were approximately equal, on average, across participants. This high variance across participants for all mapping functions was probably due to the structure of the experiment; which involved a period of random search before the vibrotactile feedback could become a source of assistance.

Skew

In addition to means and standard deviations, the distribution skew values for participant's find-time distributions were examined, and a near universal positive skew discovered; the distributions were one-tailed, with tail in the direction of increasing time. This positive skew was, again, probably due to the experiment's structure; while physical limits imposed a hard lower limit on targeting time (one can only move and respond so fast), there was no such hard upper limit. If anything, the psychological and physical need to pace one's self during a trial meant that the longer a trial lasted, the longer it would continue to last.

To examine how find-time distribution skew varied across mapping functions, the distribution skew for each of the four mapping functions was divided by the maximum of the distribution skew values across each of the four mapping functions for each participant. Next the resulting normalized skew values for each of the four mapping functions were averaged across participants and graphed (see Figure 35). This process revealed that the Perceived Gaussian and Natural find-time distributions were most skewed, while the Threshold and None distributions were least skewed. From one perspective, this result seemed counterintuitive; since the None and

Threshold functions presumably entailed random search to a greater degree than the Natural and Gaussian functions, one might expect that the distributions for the None and Threshold functions would be closer to standard normal. From another perspective, however, the result makes sense. When participants were having trouble finding the target through Natural and Perceived Gaussian mapping functions, they at least had the assistance of the gradient to guide them. When participants were having trouble finding the target through the Threshold and None functions on the other hand, they were left entirely in the "in the dark". This difference might account for the observed variations in distribution tail length across the mapping functions.

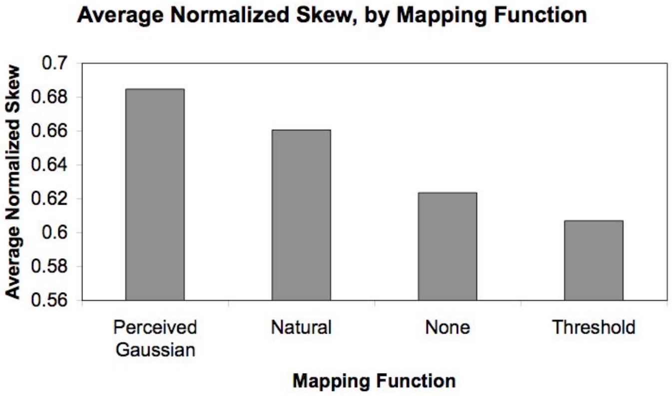

Figure 35: The average normalized skew of participants' find-time distributions, by mapping function. (Skew was normalized across mapping functions by the maximum skew value of all four functions, then averaged across participants). Average distribution skew was found to be positive for all mapping functions; the distributions were one-tailed with tail in the direction of increasing time. The degree of skew varied from mapping function to mapping function; the Perceived Gaussian distribution was the most skewed, while the Threshold distribution was the least skewed.

The fact that find-time distributions were skewed had implications for the statistical tools that could be appropriately applied during analysis. Since a standard normal distribution could not be assumed, statistical comparisons would have to be made using non-parametric tests. (In particular, the Wilcoxon Signed Rank Test could be applied, since it compares one or more subjects on a continuous measure, such as time, under two different sets of conditions.) Distribution skew also informed the choice to rely on median values rather than means as a chief measure of central tendency.

Since the experiments had not been conducted before and it was unclear what sorts of data or trends might emerge, I opted not to discard outliers from either experiment, nor to remove any of the participants from the analyses.

5.3.2.4. Choice of Mapping Function

In analyzing the relationship between find-time and mapping function, the goal was an ordering of the four functions from "quickest" to "slowest" that accurately described most (if not all) participants' trial data from the first experiment. It was hoped that the process of constructing this ranking would clarify the relative merits of each function, and indicate one "fastest" function.

A first pass at ordering the mapping functions was based solely on differences between median find-times, without consideration for statistical significance. Median find-times were found for each participant for each mapping function and placed in order (from least to greatest) for each participant. The frequency of the various orderings of the mapping functions across participants was then tabulated, and the results graphed (see Figure 36).

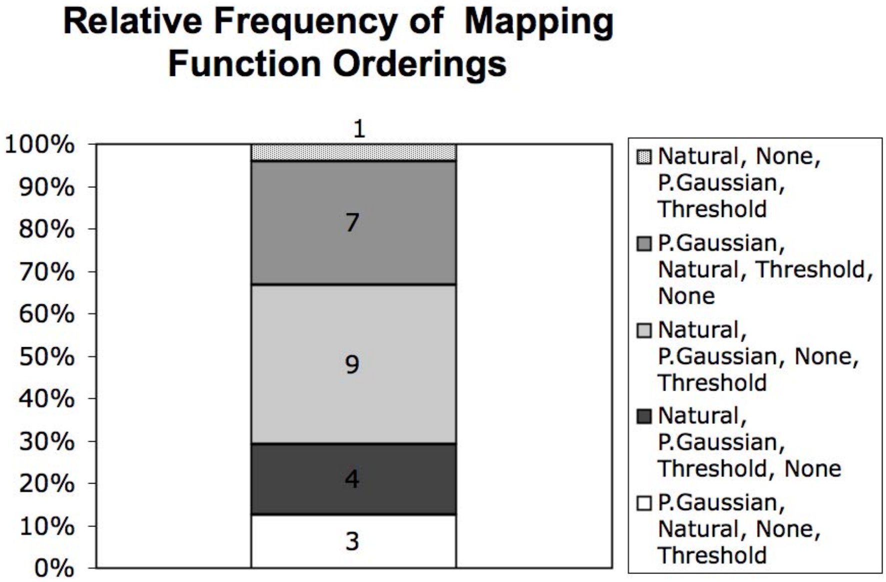

Figure 36: Relative frequency of the mapping function orderings. (The orderings are from quickest to slowest). The [Natural, Perceived Gaussian, None, Threshold] and [Perceived Gaussian, Natural, Threshold, None] rankings were most common at (9/24) and (7/24), respectively, while the [Natural, None, Perceived Gaussian, Threshold] and [Perceived Gaussian, Natural, None, Threshold] rankings were least common, at (1/24) and (3/24), respectively. (The number on or above each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

Five unique orderings appeared in participants' data; many more than the desired 1, but many less than the possible 24. The two most common orderings were [Natural, P.Gaussian, None, Threshold] (9/24) and [P.Gaussian, Natural, Threshold, None] (7/24). As anticipated, the Threshold function was not the most frequent "fastest" function. In contrast with expectations, the Perceived Gaussian mapping function was also not the most frequent fastest mapping function; though P. Gaussian was fastest for 10 of the 24 participants, Natural was the fastest function for 14 of the 24 participants. The Threshold function appeared to be less helpful for targeting than the None function (the control condition); Threshold was most frequently (13/14) the slowest function, while None was the slowest function with slightly less frequency (11/14).

This first pass based solely on median find-time values for each function provided a rough sense of an ordering: Natural and P. Gaussian at the two fastest positions (23/24), and None and Threshold at the two slowest positions (23/24). This left four possible permutations. The advantages of any one permutation over any of the others remained unclear, and the statistical significance of all rankings remained unknown.

To further clarify and solidify an ordering, the data was examined statistically at a confidence interval of 95%. Using the Wilcoxon Signed Rank test, find-time distributions were compared against each other in a pair-wise fashion to assess whether one function's find-time distribution was significantly greater-than, less-than or equivalent to the other. (In these tests, less-than corresponded to faster, while greater-than corresponded to slower). The frequencies of each of the three possible outcomes for each of the six comparisons were tabulated across participants, then graphed (see Figure 37).

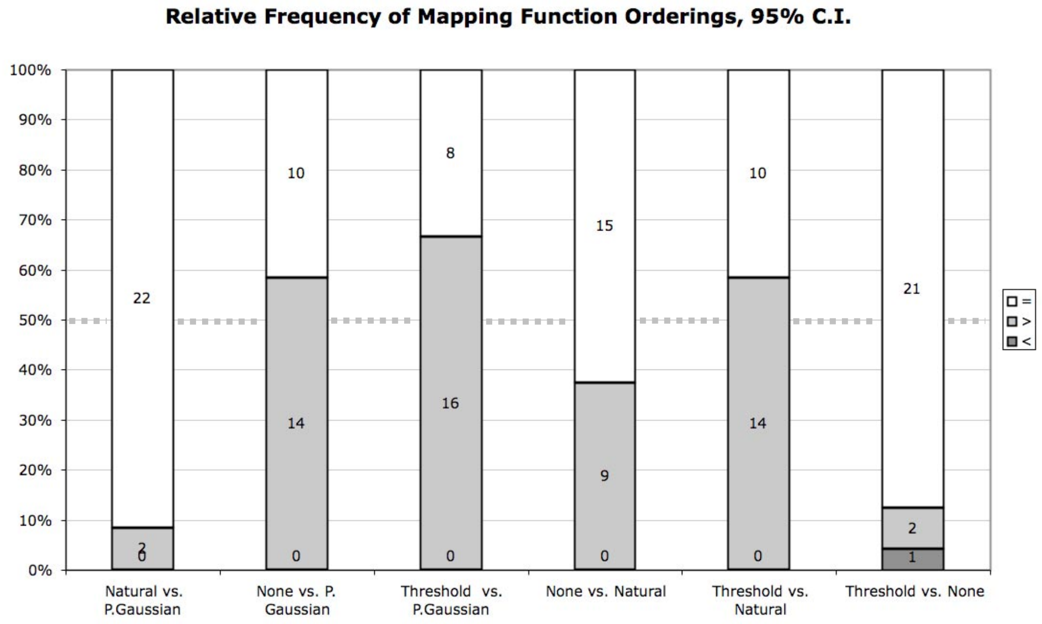

Figure 37: Pair-wise comparisons between the find-time distributions of each mapping function for all participants. Light grey, dark grey and white indicate the relative frequency of slower-than, faster-than and equal-to results for each pair-wise comparison, at a confidence interval of 95%. (The number on or above each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

In this graph, the white bar-sections represent the frequency with which no statistical difference was found between the find-time distributions for the first and second function (see x-axis labels) across participants. Light grey indicates the frequency with which the first function was significantly slower (longer find-times) than the second function, and dark grey indicates the frequency with which the first function was significantly faster (shorter find-times) than the second function. At a confidence interval of 95%, all six comparisons resulted in ties frequently: between 30% and 90% of participants' comparisons between any given pair of mapping functions were ties. The middle four comparisons (None vs. P. Gaussian & Natural; Threshold vs. P. Gaussian & Natural) most decisively suggest differences between the mapping function pairs; their non-ties are unanimously greater-thans, and their tie frequency is far less than the frequency of ties produced by the two outer comparisons (P. Gaussian vs. Natural and Threshold vs. None).

The left-most comparison indicates that Perceived Gaussian was significantly faster than Natural for two of the 24 participants. A slight bias in favor of Perceived

Gaussian as the fastest function is also indirectly evident from the Gaussian vs. None and Natural vs. None comparisons, as well as the Threshold vs. Gaussian and Threshold vs. Natural comparisons; in these comparisons, Gaussian was the faster function with greater frequency than the Natural function. Similar direct and indirect comparisons between the Threshold and None mapping functions reinforce the earlier observation that the Threshold function seems most frequently to be the slowest mapping function.

Aside from slightly tipping the balance from Natural to Gaussian as the most frequent fastest function and reinforcing Threshold's position overall as the slowest function, examining the data through pair-wise Wilcoxon tests at a confidence interval of 95% did not suggest anything fundamentally new about the ordering of the mapping functions; the most prominent feature of the graph is the large percentage of ties in all comparisons.

Since the aim of the analysis was to obtain an ordering that seemed best for the majority of participants, The pair-wise comparisons were repeated at the lower confidence interval of 80% to see how this might affect the ratio of ties to non-ties for each comparison. The results of these comparisons appear in Figure 38.

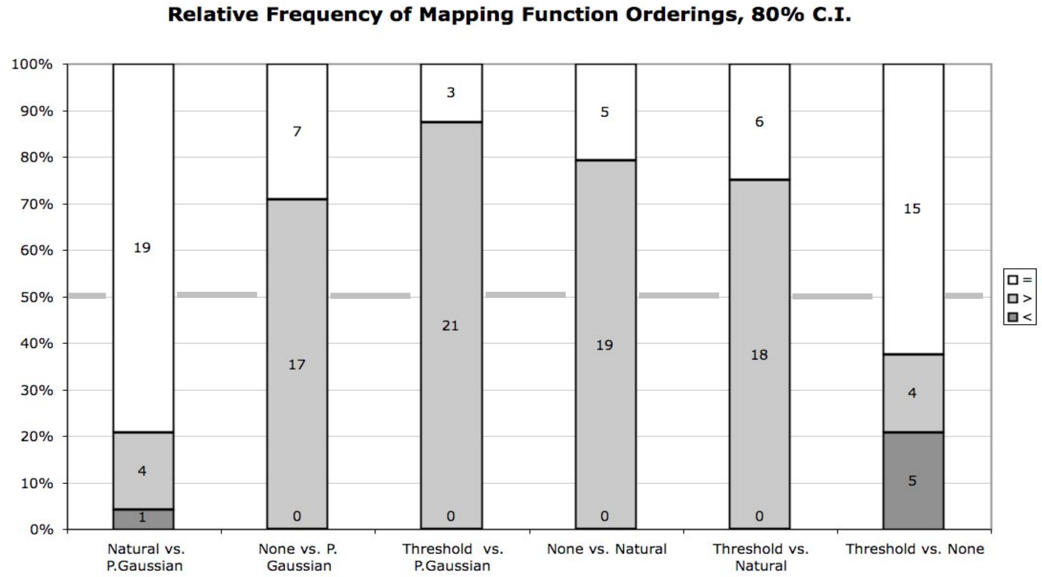

Figure 38: Pair-wise comparisons between the find-time distributions of each mapping function for all participants. Light grey, dark grey and white indicate the relative frequency of slower-than, faster-than and equal-to results for each pair-wise comparison, at a confidence interval of 80%. (The number on or above each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

With the confidence interval reduced to 80%, ties become much less frequent. For 4 out of the 6 comparisons, ties appear with a frequency of less than 30% and non-tie results became the majority. With this lowering of the confidence interval, the ratio between less-thans and greater-thans for any given comparison does not change significantly; this suggests that the trends might have been statistically significant with greater frequency at a more acceptable confidence interval had the experiment achieved greater control, or had larger sample sizes been used.

From pair-wise statistical comparison, it appears that the single ordering of mapping functions that best explains participants' time trial data is, from fastest to slowest, [P. Gaussian, Natural, None, Threshold]. Within this ordering, however, differences between P. Gaussian and Natural and between None and Threshold are slim.

To see if the differences between the mapping functions could be teased apart further, each participant's data set was divided by the experiment's two values for mapping-range (18cm and 32cm), and both halves separately examined. Splitting the data in this way seemed promising, since it was presumed that the advantages of the mapping functions would be more evident for the larger value of the mapping function than for the smaller value. A quick numerical check seemed to support this presumption; with participants' data sets split, the ratio between the fastest and slowest function's median find-times was larger for the larger value of mapping-range (~0.7 as opposed to ~0.5), and this suggested that the mapping functions' median find-times would be more "spread out" for this larger value. (These ratios were obtained by normalizing the maximum and minimum median find times across mapping functions at each range for each participant, then averaging them across participants at each range.)

When participants' data sets were split by the two values for mapping range and the relative frequency of mapping function's median find-time orderings tabulated and graphed (see Figure 39), there were in fact many more orderings for the mapping-range of 18cm than for the mapping-range of 32cm. This suggests that the guiding vibrotactile stimulus had a greater effect for the larger value of mapping-range than for the smaller one.

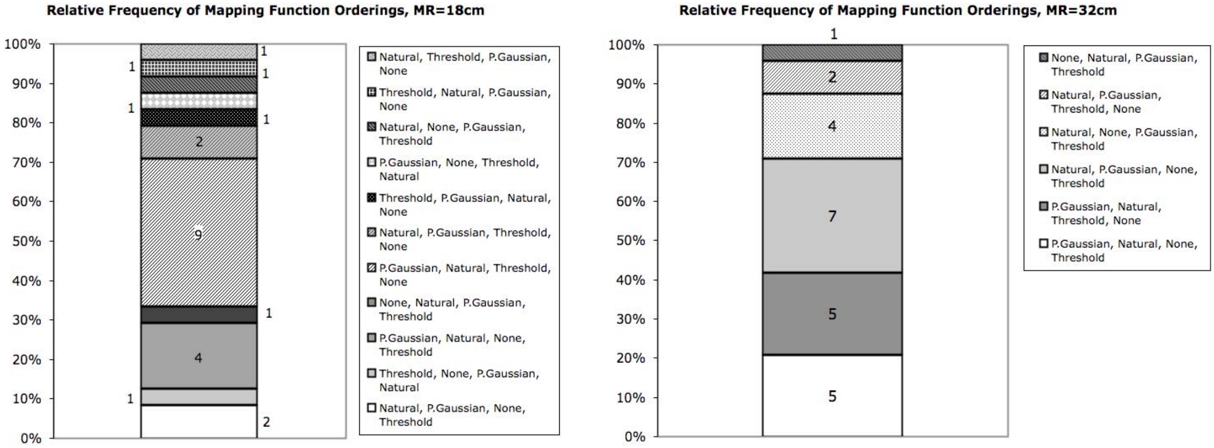

Figure 39: Relative frequency of the mapping function orderings, for two different values of mapping-range: 18cm and 32cm. (The orderings are from quickest to slowest). The number of unique orderings is greater for the mapping-range of 18cm than for the mapping-range of 32cm. Comparison of the two charts reveals that orderings may have some dependency on mapping-range. Natural is most frequently the fastest function for a mapping-range of 32cm, but Perceived Gaussian is most frequently fastest for a mapping-range of 18cm. Threshold is most frequently the slowest function for the mapping-range of 32cm, however None is more frequently slowest for the mapping-range of 18cm. (The number on or above or next to each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

I had predicted that mapping range would not influence which mapping function was the fastest and which was the slowest, however there did seem to be a correlation; at the mapping-range of 18cm, the most frequent fastest function was Perceived Gaussian (14/24), at the mapping range of 32cm, the most frequent fastest function was Natural (13/24). The most frequent slowest function at the mapping range of 18cm was None (14/24) while the most frequent slowest function at 32cm was Threshold (17/24). Although the reversal between Natural and Gaussian at the fastest position in the ordering across the two mapping-ranges was not anticipated, it makes sense when one considers the compression of the functions in space. As values for mapping range decrease, the Natural function looks increasingly like the (slow) None function; at larger and larger ranges, the Gaussian function looks increasingly like the (slow) Threshold function. The reversal between Threshold and Step at the fastest position in the ordering across the two mapping ranges can be understood in light of participant's reported dislike of the Threshold function; its constant in-range buzz would have been more tiring, numbing and annoying for the larger of the two range values, since this larger range necessitated searching a larger in-range space.

Splitting the data by the two mapping ranges did clarify the most frequent mapping function orderings in a *relative* manner (i.e. there were fewer orderings for the mapping-range of 32 than for the mapping-range of 18), however this approach did *not* clarify an overall most-frequent ordering. There were fewer orderings (5, see Figure 36) for the data that included both values of mapping-range than there were for either of the split data sets (11 orderings for mapping-range=18cm and 6 orderings for mapping-range=32cm). Furthermore, when the statistical significance of the orderings for the split data was examined in pair-wise fashion at confidence intervals of 95% and 75% (see Appendix P), the frequency of ties for each function comparison was significantly greater than the frequency of ties had been for the data that combined both mapping ranges. The split data sets did not yield more conclusive results than the combined data. Although splitting the data sets isolated experimental variables effectively, it also reduced the number of samples by half for any given participant and condition (from N=24 to N=12), and made trends harder to identify with or without statistical tools.

After the data had been split and the results examined, the full data sets were examined once more. This time, data was grouped by two curve-types: Discrete and Continuous. (None and Threshold functions were categorized as Discrete, Natural and Perceived Gaussian functions were categorized as Continuous). At this level of granularity, the most frequent two-way ordering was clear. When the frequencies of two-way orderings (based on median find-times) was tabulated and graphed (see Figure 40), over 90% of participants' orderings placed the continuous functions in the faster spot.

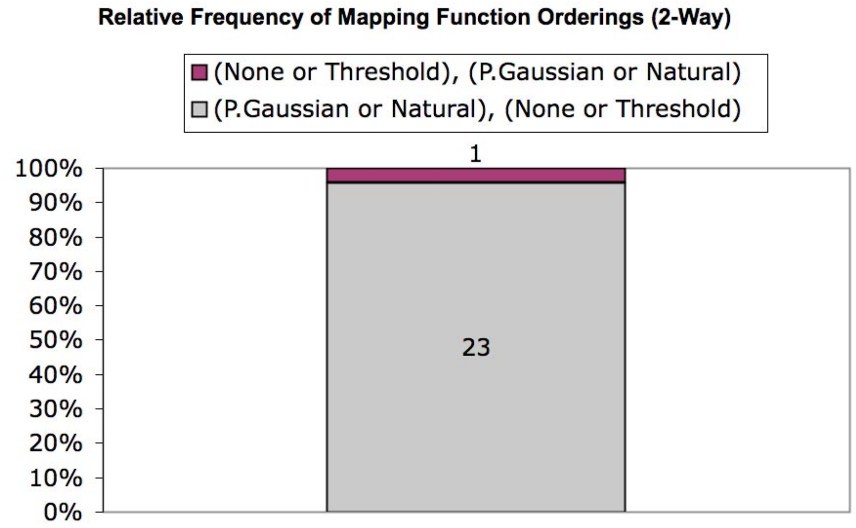

Figure 40: Relative frequency of two-way mapping function orderings. (The orderings are from quickest to slowest.) The [Natural or Perceived Gaussian, None or Threshold] ordering was nearly universal. (The number on or above each bar-region denotes the number of occurrences of the conditions represented by that bar-region).

When the statistical significance of differences was taken into account at confidence intervals of 95% and 80%, at least 80% of participants' orderings placed Continuous functions in the faster spot for a confidence interval of 95%, and greater than 95% of participants' orderings placed Continuous functions in the faster spot for a confidence interval of 80% (see Figure 41).

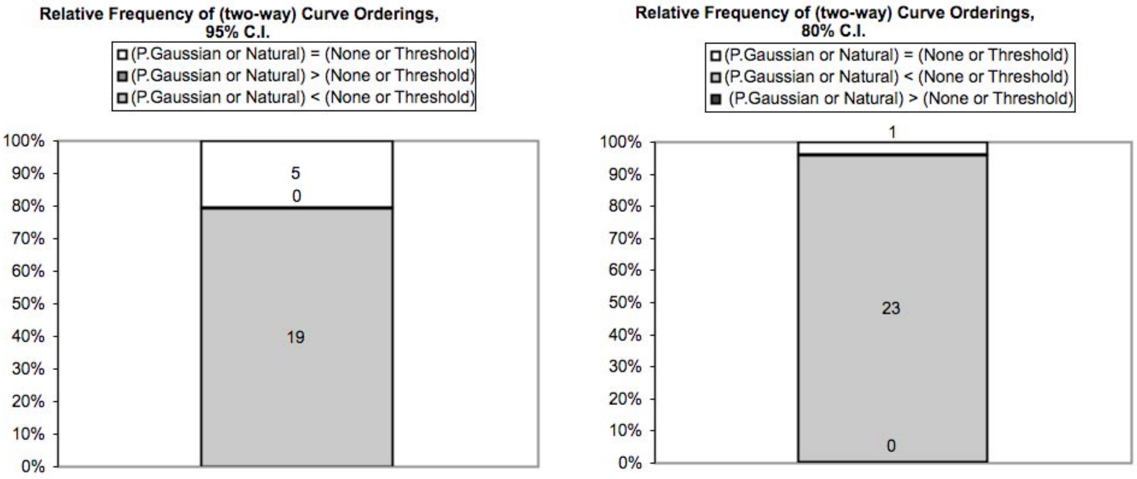

Figure 41: Relative frequency of two-way mapping function orderings. The orderings are from quickest mapping function to slowest. The [Natural or Perceived Gaussian, None or Threshold] ordering was most common.

After examination of participants' time trial data based first simply on median find-time values, then through statistical comparisons at two confidence intervals for two values of mapping range (combined and separated), it appeared that the four-way ordering of mapping functions that best explains participants' data was, from fastest to slowest, [P. Gaussian, Natural, None, Threshold]. Differences between the two Continuous functions and two Discrete functions are not stark (and may in fact vary with mapping range, though this requires further examination). The two-way ordering of mapping functions that best describes participants data is [Continuous, Discrete]. Both Continuous mapping functions clearly correlated with faster targeting, on average, than either of the Discrete mapping functions.

The second experiment required choosing one "best" mapping function, and Perceived Gaussian was chosen based on its slightly more frequent ranking as the fastest function in the statistical analysis at 95% that included data for both mapping-ranges. (This choice between Perceived Gaussian and Natural could have gone the other way if misses/cancels had been considered along with find-time distributions; though Natural and Perceived Gaussian functions together accounted for few (14%) of the missed/canceled trials, the Perceived Gaussian function accompanied approximately twice as many misses/cancels as the Natural function).

5.3.2.5. Target Size and Mapping Range

After Perceived Gaussian had been chosen as the fastest mapping function, the investigation turned toward the question of whether or not the ratio between target-size and mapping-range correlated with targeting time (for the Perceived Gaussian function). Would any one ratio correspond with shortest find-times, on average, for each participant? A predominantly visual method was used to explore the nature of the find-time/ratio relationship (though not its extent). For each participant, mean find-time was graphed as a function of the ratio between target-size and mapping-range, using the data from second experiment. Means were used rather than medians in this analysis because the available statistical analysis package, SPSS, could graph means and confidence intervals together, and this combination provided a clearer picture of distributions and central tendency than median values graphed alone). Representative graphs for three participants' data appear in Figure 42.

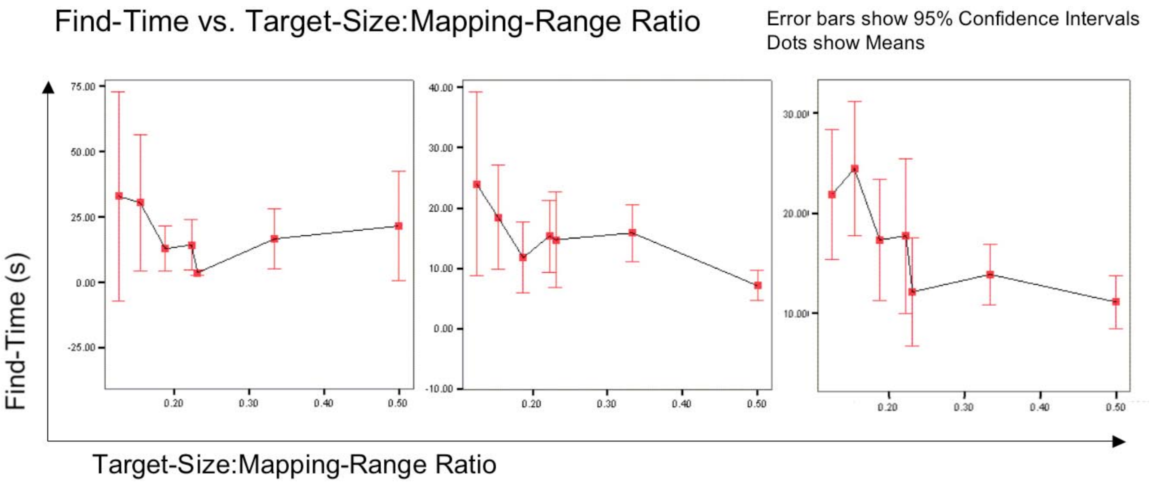

Figure 42: Sample plots of Find-Time as a function of the ratio of Target-Size to Mapping-Range, for three participants.

After the ratio of target-size to mapping-range was plotted against find-time for all participants, the graphs were classified visually as most resembling one or another of the four classification curves illustrated in Figure 43. (These curves were drawn after a quick review of participant's graphs).

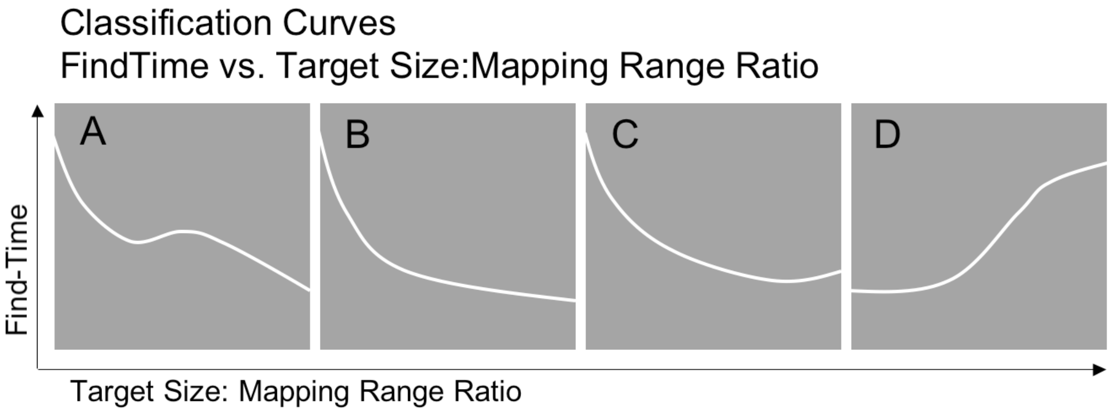

Figure 43: Curves used to visually classify participants' mean find-time as a function of the ratio between target size and mapping-function range.

This visual categorization revealed that the majority of participants' ratio graphs (15/22) resembled curve A. The remaining graphs resembled curve B (4/22), curve C (2/22) and curve D (1/22). While a central global minimum had been expected, participants' graphs appeared to head gradually downhill in the direction of increasing ratio. A quick correlation test confirmed this observation; for 82% of participants, the best linear approximation for the relationship between find-time and ratio had a negative slope. Of the statistically significant correlations (8/22 at a confidence interval of 95%), all were negative in sign and small-to-medium in extent ($|r|<0.10$ to 0.49) (Pallant, 2001). (Note: Correlation was not an ideal tool since neither linearity nor a normal distribution could be assumed; it was chosen as a secondary supplement to visual judgment because it could be applied to the data quickly and easily).

The results do not suggest that any particular ratio between target-size and mapping range corresponds to shorter average values for find-time for the Perceived Gaussian function, at least over the domain of ratios examined. While one cannot draw conclusions for ratios outside of this domain, it is possible that a global minimum might have become visible if the ratios had extended through a wider domain. Perhaps the experiment captured the drawn relationship depicted by the inset of Figure 44, but missed the bigger picture and its global minimum.

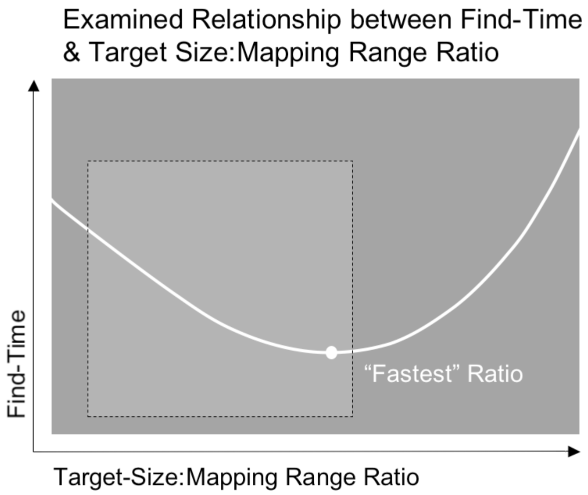

Figure 44: Expected and examined (hypothetical) relationship between find-time and the ratio of target-size to mapping-range. The experiment likely did not catch both inclines around a minimum.

The size of the search volume in relation to the mapping-range and target-size parameters may also help to explain why no clear central minimum was observed. Varying the mapping-range parameter had a strong influence on the percentage of the total search volume that fell within the region of vibrotactile response; varying the target-size parameter, on the other hand, had a comparatively weak influence on the portion of the search volume encompassed by the target. This asymmetry between the influence of changes made to target-size and the influence of changes made to mapping-range on percentages of total search volume is compounded by the natural mathematical relationship between radius and volume; both mapping-range and target-size parameters are radii, and radii become cubed in volumetric calculations. This study has focused primarily on the relationship between mapping-range and target-size; a more sophisticated study would systematically account for the fact that the ratio between target-size and mapping-range is part of a larger, more complex web of spatial relationships that includes the size of the overall search volume.

While nearly all (21/22) participants' find-time vs. ratio graphs showed a descent in find-time with increasing ratio values, most (15/22) also depicted a small "hump" near the ratio value of 0.2. Inspection of the data revealed that at this value, mapping-range fell from 32cm to 18cm, while target-size rose from 4cm to 6cm; the ratio remained at approximately 0.2, however the numerator and denominator were both changing. The consistent presence of a hump on the graphs suggests that one of the two variables in the ratio had a stronger influence on find-time than the other. It appeared that in order to understand the influence that target-size and mapping-range have on find-time, examination of a simple ratio was not sufficient; it would be necessary to consider the two variables independently.

The first independent influence considered was that of target-size. When the find-time distributions for each of the two target-size values were compared within each participants' data (through the Wilcoxon Signed-Rank test), the find-time distribution for the smaller target proved significantly smaller than the find-time distribution for the larger target, and this was the case for all participants. Clearly, bigger targets could be found more quickly. While this result seems intuitively obvious, it was perhaps especially visible from the data for two reasons: First, within the experiment, the task of holding the locator steady while pressing the locator's button was easier for larger targets than for smaller ones. Second, larger targets occupied a larger percentage of the search volume. In the design of experiments, target-size values had been chosen so as to be so small in comparison with the overall search volume that any relative differences in volume would be negligible, however the significant relationship discovered between target-size and find-time suggests that the choices made for target-size were not ideal in this respect.

After the independent influence of target-size on find-time had been examined, the independent influence of mapping-range was considered. As with examination of the find-time/ratio relationship, a predominantly visual method was employed. First, mean find-time was graphed as a function of mapping-range for all participants. For three representative graphs, see Figure 45.

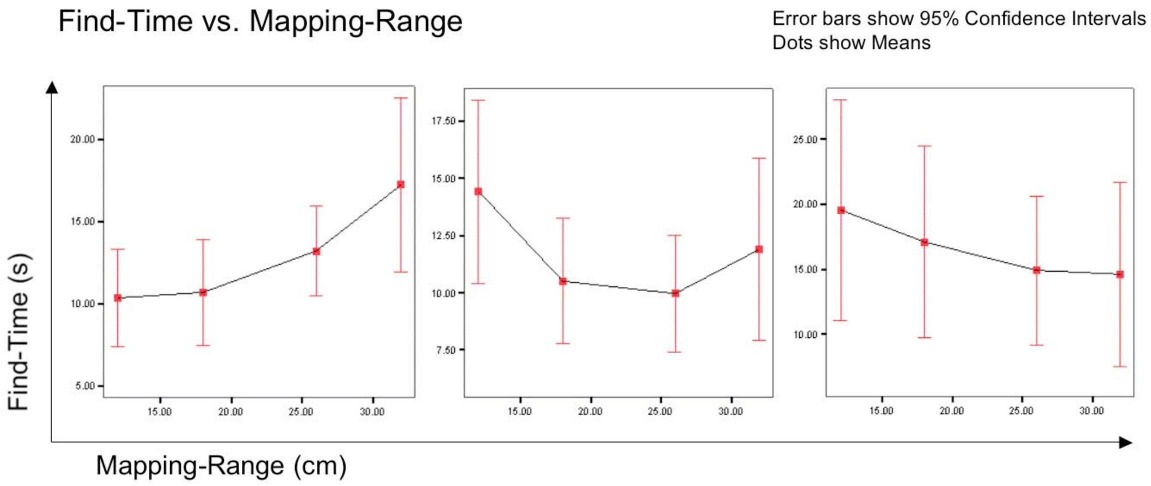

Figure 45: Sample plots of find-time as a function of mapping-range, for three different participants.

Next, the graphs were categorized according to which of the four (drawn) curves in Figure 46 they most resembled.

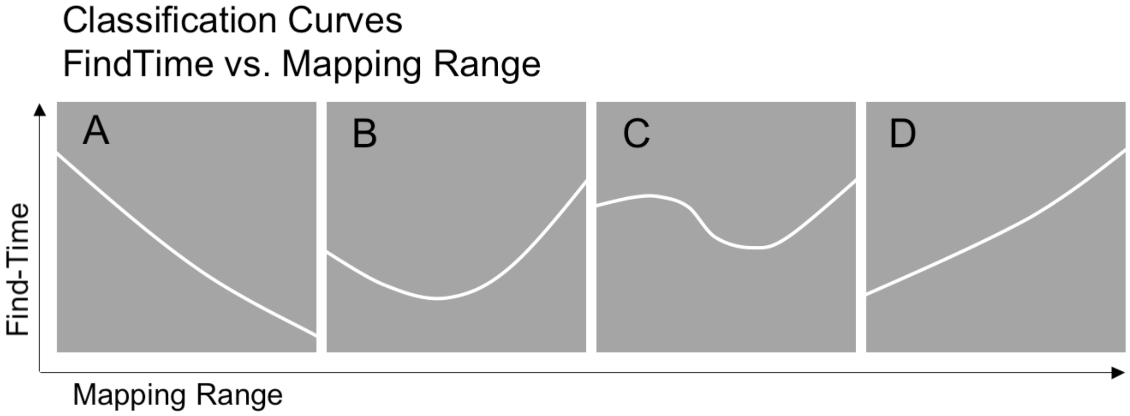

Figure 46: Curves used to visually classify participants' mean find-time as a function of mapping-range.

The results of this categorization were as follows: curve A (5/22), curve B (7/22), curve C (4/22), curve D (5/22). While I had expected that participants' find-time values might frequently fall with increasing mapping-range (at least for smaller mapping-range values), the results were evenly split across the rising and falling classification curves. In an attempt to explain this mixed result, I first considered the possibility of an interaction between mapping-range and search volume size. (As mentioned previously, changes to the mapping-range parameter altered the percentage of the search volume inside the region of vibrotactile response). Increasing the mapping-range, however, would *increase* the portion of the search volume accompanied by vibrotactile response, and thus would tend to *reduce* periods of random search preceding the onset of vibrotactile assistance. If anything, this would create a bias *toward* the expected outcome of find-time decreasing with increasing values of mapping-range[38]. Next I considered demographic differences, and found that the curve classifications did not appear to correlate with age, height, gender or handedness. It remains unclear why the data from the second experiment for each target size value did not reveal a more consistent relationship between find-time and mapping-range. Perhaps considering participants' strategies or splitting the time trial data by the two values of the target-size might reveal additional clues.

---

[38] Find-time could have been defined as the interval between entering the vibrotactile region and pressing the locator's button, rather than the interval between leaving the home position and pressing the locator's button. This alternate definition would have eliminated periods of random search at the beginning of trials (before the onset of vibrotactile assistance), but would have made time trial results brittle with respect to the possible need to enter and leave targets multiple times. (See Chapter 5, Section 1.3).

Examination of the independent influences of target-size and mapping-range on participants' average find-time revealed that find-time was most clearly dependent on the target-size variable. Thus, if there were any one ratio that might facilitate faster targeting on average, it would likely be a weighted one. The data, however, did not suggest the existence of a fastest ratio; perhaps because the domain over which the ratio was varied was too small to catch descending and ascending edges around a global minimum, perhaps because no such minimum actually existed.

5.3.2.6. Fitts' Law

After analyzing the relationship between find-time and target-size:mapping-range ratio, it seemed wise to examine the influence of Fitts' law's independent variables (target-size and home-target distance; see Chapter 2, Section 5) independently before examining their combined influence. The previous analysis had found that find-time decreased with increasing target-size, a result potentially consistent with Fitts' law, and so I chose first to focus on the relationship between find-time and target-home distance.

In order for Fitts' law to hold, average find-time would have to increase with target-home distance. To examine whether or not this was the case, mean find-time was graphed as a function of target-home distance for each participant. Data from the first and second experiments (for the Perceived Gaussian function) were combined in order to obtain the greatest possible number of target-home values, and all values for mapping-range were included in the distributions. Three of the participants' graphs appear in Figure 47.

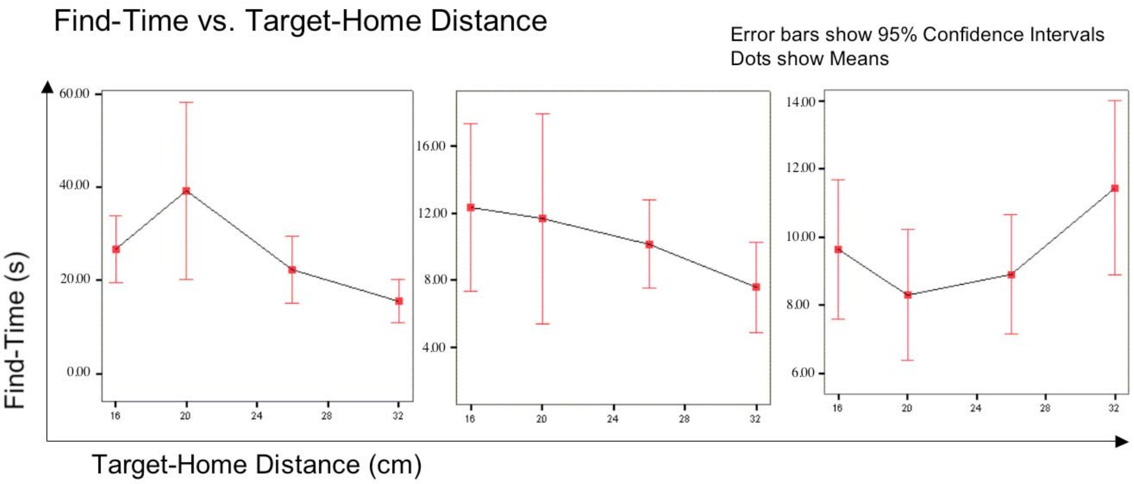

Figure 47: Sample plots of find-time as a function of target-home distance, for three different participants.

Upon visual inspection, there did not appear to be any clear pattern of find-time increasing with target-home distance. To supplement visual judgment, find-time was correlated with target-home distance. Roughly half of the correlation-slopes were found to be negative; find-time's linear approximation actually decreased with target-home distance half of the time. (Again, since neither linear relationships nor normal distributions could be assumed, linear correlation was not a particularly appropriate tool; it was merely the quickest numerical check on subjective visual judgment available.) To eliminate the possibility that combining the data from the two experiments had introduced problems, the correlations were examined for the data from both experiments independently. (Again, all values for mapping range were included in the distributions.) The signs of the correlation-slopes remained ambiguous.

Find-time did not appear to increase with target-home distance, and so it seems unlikely that Fitts' law could describe the sort of targeting that took place within these experiments. In hindsight, this apparent inapplicability is not surprising; Fitts' law describes pure targeting, while the experiments actually combined search and targeting tasks[39]. In the sorts of targeting tasks that are described well by Fitts' law,

---

[39] Find-time, for the purposes of the experiments, was defined as the time interval between leaving the home position and pressing the button within the target. If find-time had been defined differently (for example, as the interval between entering the range of vibrotactile feedback for the first time and pressing the button within the target, or as the total time spent in range before pressing the button within the target), Fitts' law might have described the data. Such alternate definitions for find-time, however, would have made time trial results more brittle in the face of the possible need to enter and leave targets multiple times. (See Chapter 5, Section 1.3).

vision typically allows a targeter to take aim at the outset of targeting; in the experiments, vision could not be employed in this fashion. The target's location became apparent gradually rather than initially. In the semi-blind vibrotactile targeting explored within the two experiments, successive approximations of manual motion were of more total importance than in traditional targeting situations where visual aim is possible. Pursuit of other mathematical relationships that might effectively describe semi-blind vibrotactile targeting is left for future work.

5.3.2.7. Learning and Fatigue

In addition to analyzing the data to answer the three research questions stated in Chapter 4, Section 6, it seemed important to consider the possible influence of learning and fatigue, since these variables clearly had the potential to confound the results. If fatigue were an issue, participants' find-times would likely grow longer as trials wore on. If learning were an issue, the reverse seemed probable. If neither fatigue nor learning were significant factors (or if the two factors were evenly balanced against each other), no relationship between trial order and find-time would be observable. (Since trials were presented in random order, it was unlikely that the various experimental conditions themselves would result in any overall relationship between find-time values and trial order).

To check for the possible combined influence of learning and fatigue, individual find-time values were plotted in the order of their respective trials for each of the participants. Initially this was done for the data from the first experiment, the experiment where learning would presumably exert the greatest influence. Visual inspection revealed no clear relationship between trial order and find-time across participants. Splitting the data for each participant by all of the experimental variables (mapping-function, mapping-range and target-home distance) then plotting find-time vs. trial order for each set of conditions similarly revealed no clear trends. When the slopes of linear regression curves superimposed on the graphs were visually tallied as 1s (positive slopes) and 0s (negative slopes), averaged within participants, and finally averaged across participants, the resulting value was 0.6; there were approximately as many positively-sloped fit lines as negatively-sloped fit lines across experimental conditions across participants. (Linear regression was not an ideal technique to use here since it presumes normal distributions of data; it was simply the most expedient way to provide a quick numerical check on subjective visual judgment). When the data from the second experiment was examined through the same procedure, no trends were visually apparent and the average slope tally across participants remained at 0.6. The fact that no clear trends were observed between find-time and trial order in the data for either experiment suggests that the combined impact of learning and fatigue was negligible.

5.3.2.8. Sources of Error

Numerous sources of error may be rectified to obtain greater experimental control; this section lists the major factors I've identified and discusses their possible impact on the results.

Software & Sensor Malfunctions At various points throughout the two experiments, the Polhemus and simulation software would temporarily or permanently hang. This meant that some trials had to be repeated, and rendered the trial times for some trials meaningless (with values on the order of several hundred or several thousand seconds or values given in negative time). Many of these were caught and the trials discounted, however a few of these with normal-looking trial times may have survived to corrupt the results for certain isolated trials.

Environmental Interruption The experiments were unavoidably situated in a somewhat busy space, and interruptions likely made some trial times longer than they would have been under more ideal conditions.

Button-Press Movement Participants' tendency to move the locator out of the target while pressing the button resulted in hits being recorded as misses. This was automatically corrected in post-processing (as discussed in Chapter 5, Section 2). Ideally, no false misses would have been recorded as such in the first place.

Motivation & Concentration Participants' motivation and concentration varied, and varied from trial to trial. While these inconsistencies were inevitable, and constitute a basic limitation of such studies, participants could have been informed more completely as to the rigid behavioral requirements of time-trial testing. This might have lent additional consistency to the results.

Facilitator's Bias Participants may have wished to "help" prove a difference between mapping functions, and this may have impacted their questionnaire responses or even the time trials themselves. A better experimental design might have involved double-blind administration of the questionnaires and time trial tests.

Combination of Random Search and Targeting Since the experiments combined targeting with some degree of random search, the results likely displayed more variation than might be expected with a pure targeting task. Including some degree of random search was unavoidable, since vibrotactile feedback was non-directional and participants had to build up a sense of target placement over time in an integrative and iterative fashion. Perhaps the extent to which random search was included could have been minimized to reduce variance in the trial time results.

Small Sample Size In an attempt to keep the experiments short while including the numerous experimental variables necessary to answer the research questions, the number of repetitions for each set of experimental conditions was lower than it might have otherwise been. This made trends hard to spot and harder to prove. For more certain results, the number of repetitions for each set of experimental conditions could have been increased (though the number of experimental variables would have had to be decreased accordingly).

Target Size As discussed in Chapter 5, Section 3.2.5, the target sizes chosen for the experiment may have been large enough that changes in the target size significantly changed the proportion of the search volume that *was* the target. Addressing this possible source of error would involve numerous tradeoffs. If the target were made smaller, pressing the button while keeping the locator in the target would become more difficult. If the search volume were made much larger, it could not be explored through arm movement alone, and would not reflect the effective read-range of present-day RFID systems.

Outliers The conservative choice not to remove outlying data points (or their respective participants) from the data during analysis may have made trends less obvious. Trends might have appeared more clearly if outliers had been removed.

Strategies Participants explored different strategies while targeting, and the strategies they reported sometimes varied with experimental conditions. This made comparisons between experimental conditions an apples-and-oranges affair, to some degree or another. While participants could have been asked to adopt one common strategy and eliminated from the analysis if they did not, it was more interesting to simply observe the various techniques participants applied than to attempt to limit this variety. Furthermore, monitoring the adoption of a common targeting strategy would be difficult to do, and might result in brittle test results, results of little meaning outside of laboratory conditions. People will always invent.

Reducing or eliminating these sources of error might strengthen the consistency and conclusiveness of results.

6. Conclusions

This chapter summarizes the experiments' conclusions, discusses their practical implications for the design and development of feedback for a hand-held locator, and mentions limitations of the study.

6.1. Experimental Conclusions

6.1.1. Influence of Mapping Function on Targeting Time

The quantitative data from the first experiment conclusively showed that the targeting times for each of the two Continuous mapping functions (Perceived Gaussian and Natural) were, on average, smaller than the targeting times for either of the two Discrete mapping functions (Threshold and None). The qualitative data revealed a strong preference among the participants for the Continuous mapping functions, and this suggests that targeting time is not an inappropriate performance metric upon which to choose a "best" mapping function.

The four-way ordering of mapping functions that best described the time-trial data for the greatest number of participants was, from fastest to slowest, [Gaussian, Natural, None, Threshold]. The overall differences between Gaussian and Natural and between None and Threshold were slight, and within the overall four-way ordering, these respective two-way orderings appeared to vary with mapping-range. This inconclusive result requires further examination.

6.1.2. Influence of Target-Size:Mapping-Range Ratio on Targeting Time

When the relationship between find-time and target-size:mapping-range ratio was graphed for all participants, no consistent global minimum appeared over the domain of ratio values examined for the majority of participants. When the respective influences of target-size and mapping-range on find-time were examined independently for all participants, find-time clearly decreased with increasing target-size but did not vary consistently with mapping-range. This lack of any consistent relationship between mapping range and find-time in the trial data from the second experiment has yet to be explained.

6.1.3. Applicability of Fitts' Law

Though targeting time was conclusively found to decrease with increasing target-size, targeting time did not appear to increase as the distance between the home position and target increased. This second result suggests that Fitts' law does not effectively describe the semi-blind vibrotactile targeting investigated in this study.

Fitts' law typically describes targeting activities for which the location of the target is visible from the outset if targeting. In this study, the target was invisible, thus targeting involved some degree of random search. The lack of visual feedback (and the bearings that visual feedback instantly provides) may explain the apparent inapplicability of Fitts' law.

6.2. Practical Implications for the Design of Feedback for a Hand-Held Locator

The experimental findings have several practical implications for the design and development of a hand-held locator.

First, the experiments demonstrate that vibrotactile gradients *can* be employed to localize targets within shelf-sized spaces. Felt vibration may not guide as sensitively as sound (as informal testing revealed – see Chapter 4, Section 3.3), however vibrotactile feedback is clearly a viable option given the goals of minimizing disruption or maintaining privacy.

Second, the experiments inform an opportunity for change at a timely moment. Present-day systems for supporting physical-digital association (e.g. bar codes, EAS and RFID) universally provide feedback to their users through a Threshold-like mapping function: the function shown in experiments to support the slowest overall targeting. As hand-held locators based on such systems (particularly RFID) begin to appear commercially in various contexts – see Davis (2004) for an early example – there will at once be an opportunity to define new behavioral conventions for a new class of appliance, and a tendency to equip the locators with the same sort of feedback their underlying implementations have always been equipped with. The experimental finding that Continuous gradients support locating faster and more acceptably than the more traditional Threshold-style feedback can inform the design of hand-held locators at a moment when the conventions for such appliances are still in flux.

Third, the fact that two seemingly very different Continuous mapping functions (Natural and Gaussian) scored so evenly in the ranking of mapping functions suggests that what matters for targeting is that the function be Continuous; the specific Continuous function employed appears to be of less importance. This implies that certain subsystems of a locator can be implemented effectively with very low precision. For instance, if a locator-item distance estimate is mapped to amplitude of vibration within the locator's microprocessor subsystem via a look-up table, the size of that look up table can be made so small as to not to impose memory constraints on the choice of microprocessor. If, on the other hand, the locator-item distance estimate is mapped to amplitude of vibration through a mathematical function, that function can be made quite simple to minimize calculation time. Minimization of calculation time and table memory within today's powerful microprocessors may constitute relatively minor advantages; a more important implication is that vibrotactile and distance-sensing subsystems do not need to be precise, nor precisely linear. Relatively crude proximity estimates (such as are possible with RFID) may suffice, and cheaper, less precise vibrotactile elements may serve in place of more precise but more expensive ones with little noticeable difference in performance. High-precision can be difficult, practically, to attain in sensors and effectors, and the result that the two very different Continuous functions ranked almost equivalently suggests that high precision may be unnecessary for the sensors and effectors of a vibrotactile hand-held locator.

Fourth, since the experiments did not clearly suggest any one target-size:mapping-range ratio that corresponded with quickest overall targeting, the idea of making mapping range dependent on the size of a sought item in order to facilitate targeting can probably be discounted. Discounting this apparently unsuccessful optimization simplifies the system that would need to be in place for the locator to function; locatable items would not need to be associated with digital representations of their own size.

Finally, since Fitts' law did not effectively describe participants' data, it appears that this relationship cannot effectively inform the design of interfaces that rely on a vibrotactile locator. The inapplicability of this mathematical model poses no great stumbling block, however; the reach of mathematical models in interface design is limited, and measurement-based approaches constitute a small fraction of the spectrum of design approaches that are available.

6.3. Limitations

The experimental conclusions and their possible implications must be considered along with certain caveats. In addition to the sources of experimental error (discussed in Chapter 5, Section 3.2.8) and the modeling assumptions (discussed in Chapter 4, Section 5.1), the strength of conclusions is limited by the ordering of the two experiments, the fact that the analyses focused on the nature of statistical differences in targeting time but not their extent, and the abstractness of the experimental setting and task.

In conducting the two experiments, a linear approach was adopted where the outcome of the first experiment determined values for the variables used in the second. As a result, interactions between certain variables of interest in the two experiments could not be examined. (for example, it is conceivable that Fitts' law holds when the mapping function is not Perceived Gaussian; the experimental structure did not permit investigating such a dependency.)

When the results of the first experiment were analyzed statistically, trial time distributions for the different mapping functions were compared along a relative scale of only three values: faster-than, slower-than and equal-to. This categorization captured the nature of differences between mapping functions' find-time distributions, but ignored their extent. Thus, extreme differences may have been flattened and subtler ones may have been brought into relief. In a sense, the extent of differences was indirectly revealed through examination of the frequencies of mapping function orderings across participants, however a more detailed analysis would have addressed the extent of statistical differences in a more direct fashion.

Ultimately, targeting virtual points repeatedly in a laboratory experiment is different from finding a single bottle in a pharmacy, locating a book in a library or manually drawing together a product ID scanner with a label on a check-out item. While the experiments captured certain salient aspects of targeting in shelf-sized spaces with a hand-held locator (principally aspects pertaining to scale and manner of motion), it did not (and could not) model the numerous situation-dependent aspects that might prove important. The ecological validity of the conclusions outside of the test environment is therefore suspect, and designers would do well to verify them before accepting them uncritically for any particular application.

7. Future Work

Preliminary explorations and experiments have led to a position with several degrees of freedom for further movement. At this juncture, several directions might be taken depending on ones' scientific, technical or entrepreneurial orientation. Suggestions for future work are here grouped by three areas: experimentation, implementation and application.

7.1. Experimentation

While a certain degree of control was obtained within the formal experiments, tighter control is definitely possible. Repeating the experiments (with their various sources of error corrected for) would likely strengthen the conclusiveness of the results and shed additional light on confusing outcomes. (Is the fastest mapping function dependent on mapping-range, as the first experiment suggested? Is there no strong correlation between mapping range and targeting time, as the data from the second experiment suggested? Another iteration through the experiments would likely clarify these issues.)

With the aim of minimizing disruption within the locators' environment of use, the experiments focused exclusively on feedback of a vibrotactile mode. If this constraint were relaxed, numerous other forms of feedback would become viable, however, and could be studied through similar experiments. The human ear's sensitivity and dynamic range make sound a compelling candidate for future examination, and the sensitivity with which light is perceived – together with the directed (and thus less public) nature of visual signals – might make light a good compromise. Multimodal combinations of visual, auditory and haptic feedback could also be explored.

Though a variety of formal experiments might be conducted to verify or extend the conclusions of this study, I believe additional formal experimentation would be premature. Designed vibrotactile behaviors for a hand-held locator become most meaningful only when such locators exist and can be placed within the reach of people who can make good use of them. Additional laboratory experiments will be most valuable once they can be based upon a fully functional locator that is matched to an actual, real-world locating activity.

7.2. Implementation

Simulation can only carry the design of behaviors for a hand-held locator so far; to ultimately verify their utility, a hand-held locator (one that can sense proximity) is required. The necessary technologies exist to create such a locator, and implementation is largely a matter of stringing them together. The RFID reader design explored in Chapter 4, Section 2 and its associated appendices (A, B, C and D), together with the vibrotactile subsystem discussed in Chapter 4, Section 3.2 can serve as a basis for this work.

[Note: This discussion of future work concerns the locator and its behaviors, however it must be reiterated that the locator must ultimately work within a larger system. Tagging, categorizing, specifying and locating must all be addressed before the locator itself can provide useful digital support for finding physical things.]

7.3. Application

The approach of this work has been pragmatic with respect to technological possibility but speculative with respect to human need. The constructed use scenarios that were employed for conceptual design, demonstration and experimentation proved sufficient for clarifying certain issues and moving through a first turn of a design cycle, however these fictions must ultimately give way to – or else become – actual applications in the real world if the concept of locating sought items through vibrotactile gradients provided by a hand-held locator is ultimately to be of service.

Where might automated assistance from a hand-held locator prove most useful? Chapter 5, Section 3.1.2 delineates several problem-spaces that might fruitfully be searched for strong applications; this section might serve as a point of departure for field studies.

Since the overhead of putting a locating system in place may be significant, it may prove most valuable to consider application domains where RFID labels and hand-held readers are already in use. Where are the breakdowns? Might any be addressed by the addition of continuous vibrotactile feedback?

In evaluating potential applications for gradient-based manual guidance cues, it may additionally prove fruitful to look beyond locating digitally IDed items. Locating metal on a traveler's person in an airport with a hand-held metal detector, for example, is a task that does not rely on automated digital identification, but is one that might benefit from gradient-based feedback. This potential application seems promising in light of the speed, repetition and attention requirements that characterize work at airport security checkpoints.

While the finding of *things* may prove increasingly amenable to automation, the finding of *needs* remains at least as much an art as a science. Scenarios, participants' suggestions, current work practices surrounding RFID systems and present-day locating tasks may all prove useful points of departure in the search for a valued application.

This thesis has addressed the design of gradients in vibrotactile feedback for guiding the manual movement of a hand-held locator. The underlying technical basis for such an appliance has been discussed theoretically and verified empirically, the idea of locating physical things through digital means has been situated with respect to various streams of research, and an experimental platform has been created and verified. The study has presented one complete turn of a (vibrotactile gradient) design cycle – from concept to evaluation – and articulated the limitations inherent to both process and findings. All of these measures, however, constitute just a first step. To realize convivial digital support for locating things within the physical world, extensive work in the areas of implementation, application and experimentation remains to be done.

8. References

Akamatsu, M. and I. S. MacKenzie, 1996. “Movement Characteristics Using a Mouse With Tactile and Force Feedback” International Journal of Human-Computer Studies, 45, 483-493.

BBC, 2004. “Radio tags spark privacy worries.” [BBC News online]. Available from: http://news.bbc.co.uk/2/hi/technology/3224920.stm [Accessed March 2004].

Bellis, M., 2004. “Selling the Cell Phone Part 1: The History of Cellular Phones.” [online]. About.com. Available from: http://inventors.about.com/library/weekly/aa070899.htm [Accessed 20 June 2004].

Brave, S. and A. Dahley, 1997. “inTouch: A Medium for Haptic Interpersonal Communication.” In: Human Factors in Computing Systems: Conference Proceedings, Atlanta, GA, March 22-27, 1997: ACM Press, 363-364.

Brazil, E. and Fernström, M, 2003. “Audio Information Browsing With The Sonic Browser.” In: Coordinated and Multiple Views in Exploratory Visualization: Conference Proceedings, London, July 15, 2003. IEEE Press, 26-33.

Card, S., 2002., “Information Foraging: Instructors”. [conference program online]. Available from: http://www.sigchi.org/chi2002/tut-mon.html [Accessed March 2004].

Challis, B. P. and A. D. N. Edwards, 2000. “Design Principles for Tactile Interaction.” In: Haptic Human-Computer Interaction: Conference Proceedings, Glasgow, UK, August 31 - September 1, 2000. Brewster, S. and R. Murray-Smith, eds: Springer, 18-24.

Chang, A., O'Modhrain, S., Jacob, R., Gunther, E. and H. Ishii, 2002. “ComTouch: Design of aVibrotactile Communication Device.” In: Design of Interactive Systems: Conference Proceedings, London, UK, June 25-28, 2002: ACM Press, 312-320.

Cooper, A., 1999. The Inmates are Running the Asylum. Sams.

Cormen, T., Leiserson, C. E., Rivest, R. L. and C. Stein, 1990. Introduction to Algorithms. Cambridge, MA: MIT Press.

Coupland, D., 1996. Microserfs: Regan Books.

Darken, R. and M. Zyda, 2001. "VR - The Historical Context." [online]. Available from: http://www.movesinstitute.org/~zyda/courses/1HistoryOfVR.pdf [Accessed March 2004].

Davis, J., 2004. "No More Lost Balls." WIRED, June 2004. Available from: http://www.wired.com/wired/archive/12.06/start.html?pg=10 [Accessed June 25 2004].

Dennerlin, J. T., Martin, D. B. and C. Hasser, 2000. "Force-Feedback Improves Performance for steering and Combined Steering-Targeting Tasks." In: SIGCHI Human Factors in Computing Systems: Conference Proceedings, The Hague, Netherlands, April 1-6, 2000: ACM Press, 423-429.

Dougherty, J., 2003. "Boycott against Gillette launched." [online] Available from: http://www.worldnetdaily.com/news/article.asp?ARTICLE_ID=34051 [Accessed March 2004].

Dourish, P., 2001. Where the Action Is - The Foundations of Embodied Interaction. Cambridge, MA: MIT Press.

E.P.A., 2003. "Lead in Paint, dust and Soil." [online]. Available from: http://www.epa.gov/opptintr/lead/index.html [Accessed March 2004].

Ericson, J., 2002. "RFID Changes Everything." E-Business Ecosystem [online]. Available from: http://www.line56.com/articles/default.asp?ArticleID=4025 [Accessed March 2004].

Erp, J. B. F. v., 2002. "Guidelines for the Use of Vibro-Tactile Displays in Human Computer Interaction." In: Eurohaptics: Conference Proceedings, Edinburgh, 2002: 18-22.

Erp, J. b. F. v., 2003. "Distilling the Underlying Dimensions of Tactile Melodies." In: EuroHaptics: Conference Proceedings, Dublin, Ireland, July 6-9, 2003: Trinity College and Media Lab Europe, 111-120.

Finkenzeller, K., 2002. RFID Handbook, new York: John Wiley & Son.

Finkenzeller, K., 2004. "Frequencies for RFID Systems." [online], Available from: http://www.rfid-handbook.de/rfid/frequencies.html [Accessed April 2004].

Fitts, P., 1954. "The Information Capacity of the Human Motor system in Controlling the Amplitude of Movement." Journal of Experimental Psychology, 47, 381-391.

Franzen, O., 1969. "The Dependence of Vibrotactile Threshold and Magnitude Functions on Stimulation Frequency and Signal Level – A Perceptual and Neural Comparison". Scandinavian Journal of Psychology, 10, 289-297.

Galbourne, A., 2003. “Texas Instruments HF Antenna Design Notes”. Technical Report 11-08-26-003, Texas Instruments RFID Division, September 2003.

Gamez, E. H. C., Leland, N., Shaer, O. and R. J. K. Jacob, 2003. “The TAC Paradigm: Unified Conceptual Framework to represent Tangible User Interfaces.” [online], Tufts University, Available from: www.cs.tufts.edu/~oshaer/CLIHC2003v8.ppt [Accessed February 2004].

Gaver, B., 2002. Lecture on Interaction Design at the Royal College of Art, given at the University of Limerick, 2002.

Gershenfeld, N., 1999. When Things Start to Think. New York: Henry Holt & Company.

Gunther, E., Davenport, G. and S. O'Modhrain, 2002. “Cutaneous Grooves: Composing for the Sense of Touch.” In:New Instruments for Musical Expression: Conference Proceedings, Dublin, Ireland, May 24-26, 2002.

Hasser, C., Goldenberg, A. S., Martin, K. M. and L. B.Rosenberg, 1998. In: “User Performance in a GUI Pointing Task with a Low-Cost Force-Feedback Computer Mouse.” In: American Society of Mechanical Engineers - Dynamic Systems and Control Division: Conference Proceedings, 1998: ASME, 151-157.

Heim, M., 1993. The Metaphysics of Virtuall Reality, Oxford: Oxford University Press.

Heisse, J. J., 2003. “The Future of Virtual Reality: Part Two of a Conversation with Jaron Lanier.” [online], Sun Developers Network. Available from: http://java.sun.com/features/2003/02/lanier_qa2.html [Accessed February 2004].

Hightower, J. and Borriello, G., 2001. “A Survey and Taxonomy of Location Systems for Ubiquitous Computing.” Technical Report UW-CSE 01-08-03 University of Washington WA.

Hsiao, K.Y., 1999. Fast Multi-Axis Tracking of Magnetically-Resonant Passive Tags: Method and Applications. Masters Thesis, MIT.

Intille, S., 2002. “Change Blind Display for Ubiquitous Computing Environments.” In: Ubiquitous Computing: Conference Proceedings, Gothenberg, Sweden, September, 2002: Springer-Verlag, 91-106.

Ishii, H. and Ullmer, B., 1997. “Tangible Bits: towards seamless Interfaces between People, Bits and Atoms.” In: SIGCHI Human Factors in Computing Systems: Conference Proceedings. Atlanta, GA: ACM Press, 234-241.

Kenshalo, D., 1978., “Biophysics and Psychophysics of Feeling.” In: Handbook of Perception, Vol. VIB. Academic Press, 29-49.

Klatzky, R. L. and S.J. Lederman, 2002. In: Experimental Psychology, Vol. 4. Healy, A. F. and R.W. Proctor, eds. New York: Wiley, 147-176.

Lamping, J. and R. Rao, 1996. “Visualizing Large Trees Using the Hyperbolic Browser.” In: SIGCHI Human Factors in Computing Systems: Conferenece Proceedings, Vancouver, April 13-16, 1996, ACM Press.

Langdon, P., Hwang, F., Keates, S., Clarkson, P. J. and P. Robinson, 2002. “Investigating Haptic Assistive Interfaces for Motion-Impaired Users.” In: Eurohaptics:Conference Proceedings, Edinburgh, 2002.

Lee, Y. and P. Sorrells, 2001. microID 13.56 MHz RFID System Design Guide [online], Microchip Technology. Available from: www.microchip.com/download/lit/pline/rfid/guides/21299d.pdf [Accessed February 2004].

Lefferts, D., 2002. Interview on Use of Bar Codes in Guinardi’s Supermarkets, PA, 2002.

Leydon, K., 2003. “Sensing the Position and Orientation of Hand-Held Objects.”, Technical Report UL-CSIS-03-02, University of Limerick, Ireland.

Lloyd, G. and N. Sivin, 2002. The Way and the Word: Science and Medicine in Early Greece, Yale University Press.

Long, J. B., 2000. “History of PDA's” [online], The Webmasters, Available from: http://cctr.umkc.edu/~jblong/hist.htm [Accessed February 2004].

Mackay, W., 2000. “Augmented Reality: Dangerous Liasons or Best of Both Worlds?” In: Design of Augmented Reality Environments: Conference Proceedings, Elsinore, Denmark, April 12-14, 2000: ACM Press.

Mackay, W. E., 1996. “Augmented reality: A new paradigm for interacting with computers.” La Recherche, 284, March 1996.

MacLean, K. E. and J.B. Roderick, 1999. “Alladin: Exploring Language with a Haptic Doorknob.” In: 12th Annual ACM symposium on User Interface Software and Technology: Proceedings Asheville, NC, 1999: ACM Press.

Manzini, E., 1991. “The End of the Mechanical Age.” In: The History of Industrial Design: The Dominion of Design, Vol. 2, Pirovano, C. ed: Electa Milan, 34-55.

McCullagh, D. 2003. “RFID tags: Big Brother in small packages.” News.com, [online] Available at: http://news.com.com/2010-1069-980325.html [Accessed March 2004].

McCullough, M., 1996. Abstracting Craft, Cambridge: MIT Press.

McLuhan, M., 1964. Understanding Media, New York: McGraw-Hill.

Murray, A., Klatzky, R. and P.K. Khosla, 1998. “Enhancing Subjective Sensitivity to Vibrotactile Stimuli.” In: American Society of Mechanical Engineers - Dynamic Systems and Control Division: Conference Proceedings, 1998: ASME, 157-162.

Murray, A., Klatzky, R. and P.K. Khosla, 2003. “Psychophysical Characterization and Testbed Validation of a Wearable Vibrotactile Glove for Telemanipulation.” Presence, 12, 156-182.

Negroponte, N., 1995. “Being Decimal.” WIRED, 3(11), November 1995.

N.I.E.H.S., 2001. “Dioxin Research at the National Institute of Environmental Health Sciences.” National Institute of Health, Factsheet [online], Available from: http://www.niehs.nih.gov/oc/factsheets/dioxin.htm [accessed June 20, 2004].

Novint Technologies, 1999. “Undersea Exploration” Novint Technologies [online], Available from: http://www.novint.com/Telerobotics.htm [Accessed March 10 2004].

Pallant, J., 2001. SPSS Survival Manual, Buckingham: Open University Press.

Pang, A. S.K., 2002. “The Making of the Mouse.” American Heritage of Invention and Technology, 17(3), Winter 2002.

Paradiso, J. A., Hsaio, K., Lifton, J. and A. Adler, 2000. “Swept-Frequency, Magnetically-Coupled Resonant Tags for Realtime, Continuous, Multiparameter Control.” IBM Systems Journal, 39, 892-914.

Pasquero, J. and V. Hayward, 2003. “STReSS: A Practical Tactile Display System with One Millimeter Spatial Resolution and 700 Hz Refresh Rate.” In: EuroHaptics: Conference Proceedings. Dublin, Ireland. Trinity College and Media Lab Europe.

Petroski, H., 2000. The Pencil - A History of Design and Circumstance. New York: Alfred A. Knopf.

Power, N., 2002. Personal Communication, Limerick.

Rosenblith, L. J. ed, 2004. “Fact Sheet - Electronic Waste.” In: Environmental Issues, About.com [online] Available from: http://environment.about.com/cs/ewaste/a/ewaste.htm [Accessed March 2004].

Sellen, A. J. and R. H. R. Harper, 2002. The Myth of the Paperless Office. Cambridge: MIT Press.

Sepulveda, P., 2004. “Digital Shelters.” [online thesis website], Royal College of Art, London. Available at: http://www.crd.rca.ac.uk/phd/pedro/projects.html [Accessed March 2004].

Sharpe, B., 2001. “Information Appliances.” [online], Appliance Studio. Available at: http://www.appliancestudio.com/home/applianceness.htm [Accessed June 20 2004].

Sokal, N., 2000. “Class-E High Switching-Mode High-Efficiency Tuned RF Microwave Power Amplifier: Improved Design Equations.” In: IEEE International Microwave Symposium: Proceedings. Baltimore, June 2000: IEEE Press.

Sorrells, P., 2000. “Optimizing Read Range of RFID Systems.” EDN Magazine [online], December 7, 2000.

Stone, R. J., 2000. “Haptic Feedback: A Brief History from Telepresence to Virtual Reality.” In: Haptic Human-Computer Interaction: Conference Proceedings. Glasgow, UK, August 31 - September 1, 2000. Brewster, S. and R. Murray-Smith eds.: Springer, 1-16.

Sullivan, J., 2002. “Interactive Media Design Review: The Crossing Project.” International Design, November 2002, 58-59.

Temiskaming, 2002. “PCB Incineration: A Risk to Community Health and the Environment.” [online]. Public Concern Temiskaming, Ontario. Available at: http://members.fortunecity.com/toxic/library/pcb_incineration.htm [Accessed March 2004].

Texas Instruments, 1995. “Team of Three Texas Instruments Engineers Invented the Portable Hand-Held Electronic Calculator.” [online], Texas Instruments. Available at: http://www.ti.com/corp/docs/company/history/calcbackg.shtml [Accessed February 2004].

Texas Instruments, 2001. “Tag-It HF-I Transponder Inlays.” Product Bulletin.

Texas Instruments, 2002. “Tag-it FAQ.” [online], Texas Instruments. Available from: http://www.ti.com/tiris/docs/customerService/faq.shtml [Accessed March 2004].

Thakara, J., 2003. “The Post-Spectacular City.” In the Bubble [online], Doors of Perception, Westergasfabriek. Available from: http://www.doorsofperception.com/In+the+Bubble/ [Accessed March 2004].

Toffler, A. 1970. Future Shock. New York: Random House.

Verplank, B., Gurevich, M. and M. Mathews, 2002. In: New Instruments for Musical Expression: Conference Proceedings. Dublin, Ireland, May 24-26 2002: 33-36.

Verrillo, R. T. and G.A. Gescheider, 1992. "Perception via the Sense of Touch." In: Tactile Aids for the Haring Impaired, Summers ed., London: Whurr Publishers Ltd.

Wallington, P., 2004. (paulwallington@btconnect.com), 4 March 2004. Re: RFID Tag Pricing. E-Mail to K. Leydon (Krispin.Leydon@ul.ie).

Weinberger, D., 2003. "Why Matter Matters." [online], HyperOrg.com. Available from: http://www.hyperorg.com/blogger/mtarchive/001686.html [Accessed February 2003].

Weiser, M., 1991. "The Computer For the Twenty-First Century." Scientific American, September 1991, 94-110.

Weiser, M. and J.S. Brown, 1997. In: Beyond Calculation: the Next Fifty Years. New York: Copernicus, 75-85.

Wellner, P., 1991. "The DigitalDesk Calculator: Tangible Manipulation on a Desk top Display." In: UIST: Conference Proceedings. Hilton Head, South Carolina, November 11-13, 1991: 27-33.

Wellner, P., 1993. "Interacting with Paper on the Digital Desk." Communications of the ACM, 36, ACM Press: 86-96.

Wilson, B., 2000. "Lucid Dreaming: The Truth in Materials." [online], Chronogram. Available from: http://www.chronogram.com/backIssues/2000/04april/lucid.htm [Accessed June 20, 2004].

Winograd, T. ed., 1996. Bringing Design to Software. New York: ACM Press.

Winograd, T. and F. Flores, 1986. Understanding Computers and Cognition. Norwood: Ablex Publishing Corporation.

Appendix A. RFID Reader Schematic (Reference Design)

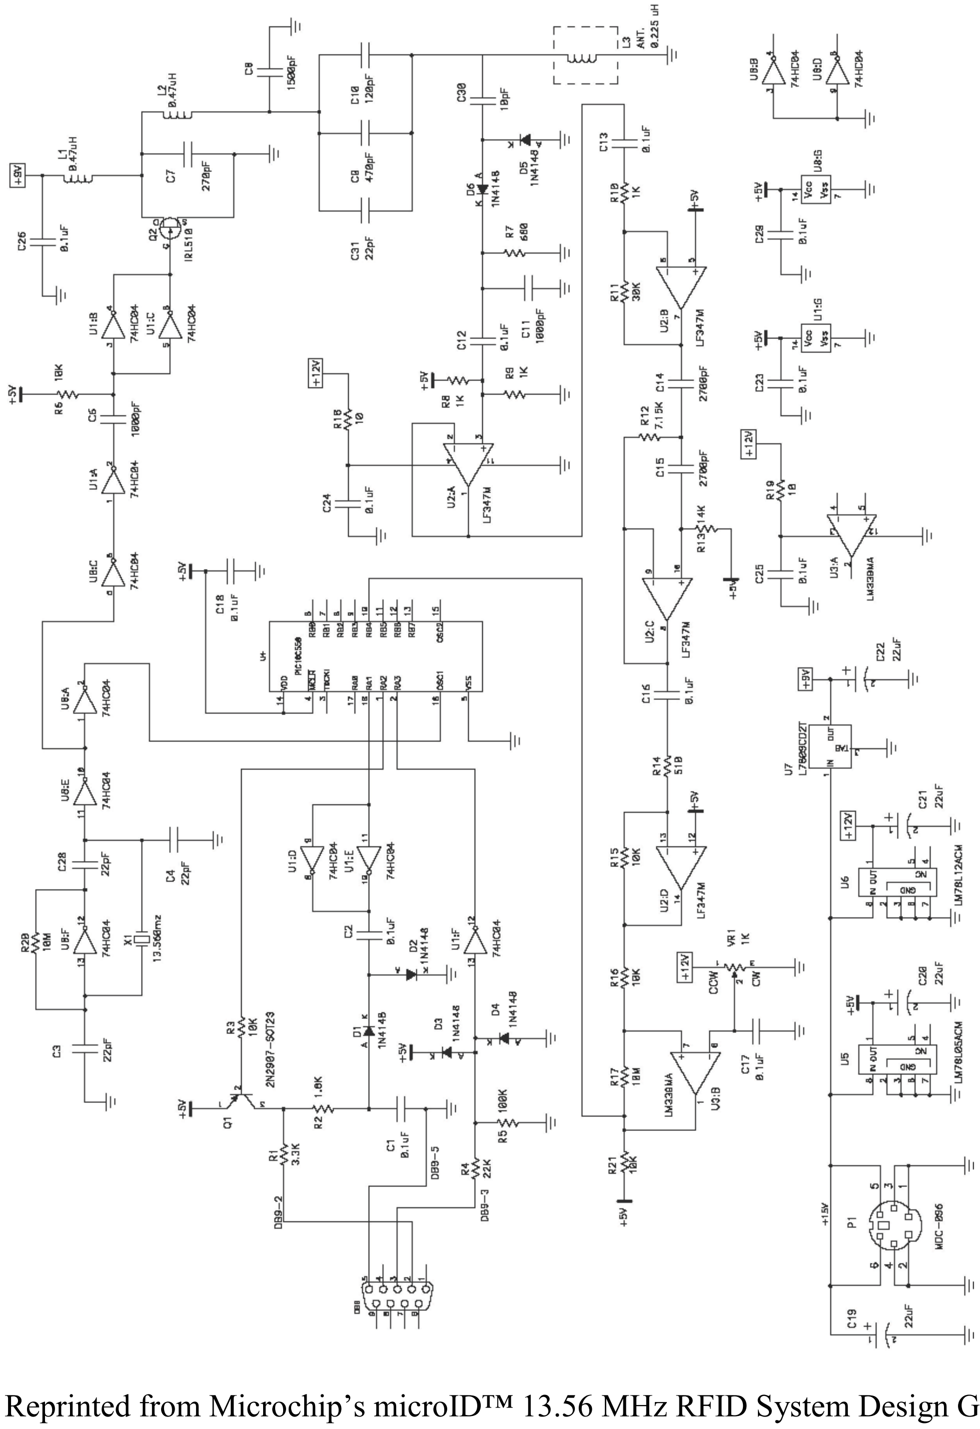

Reprinted from Microchip's microID™ 13.56 MHz RFID System Design Guide, page 103 with permission from Microchip Technologies.

Appendix B. Modified RFID Reader Design Notes

Receiver

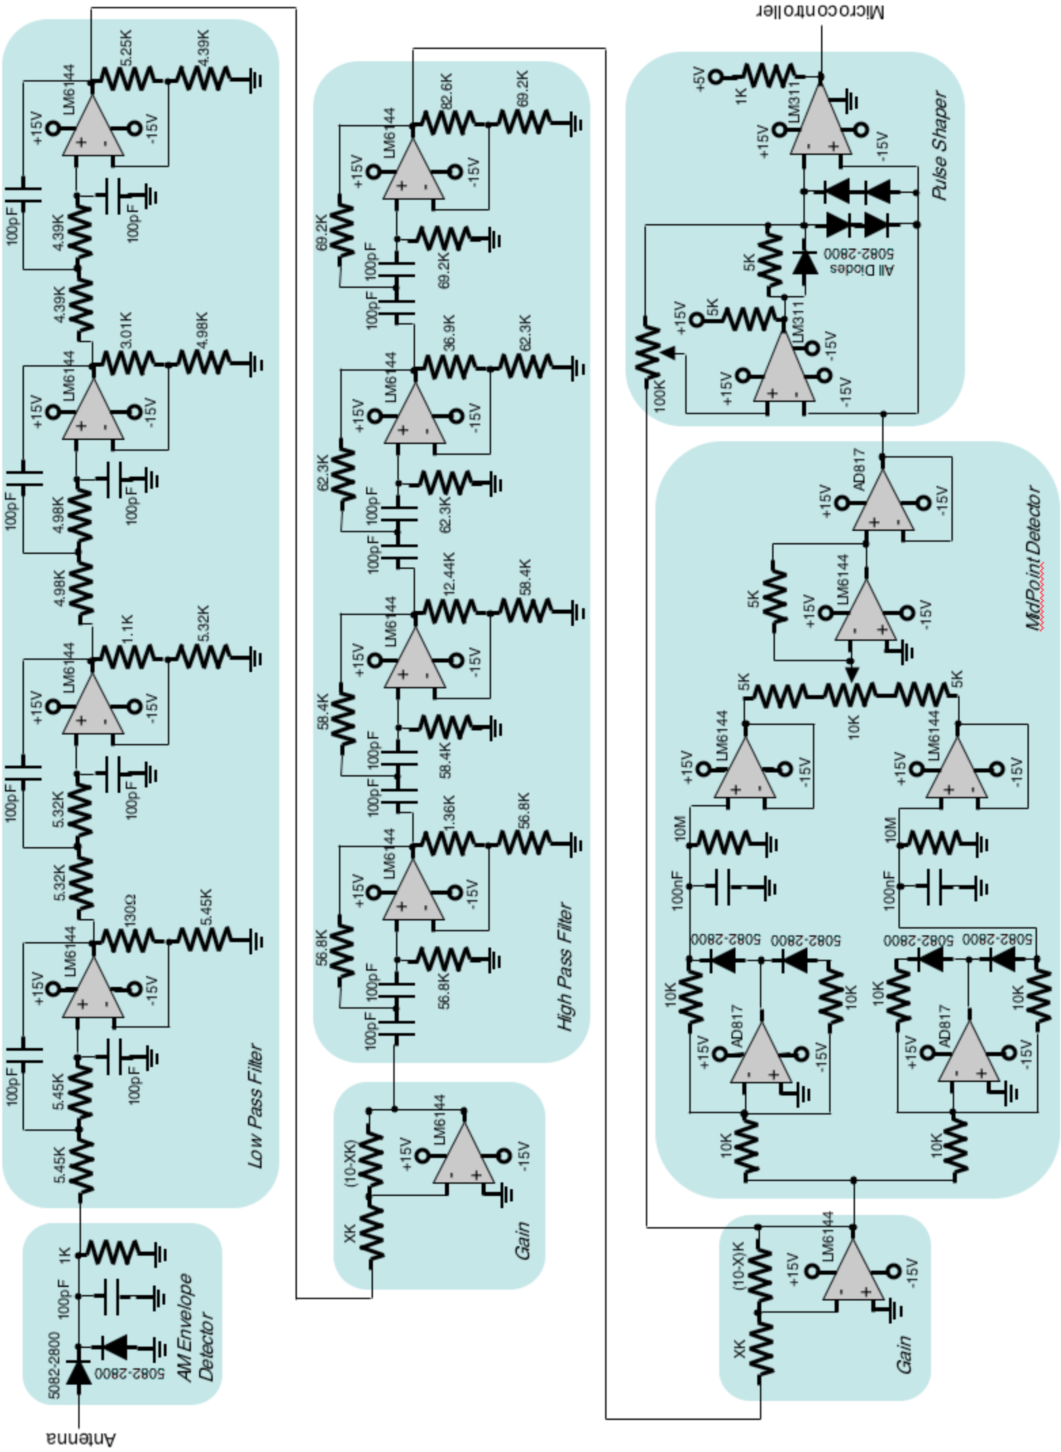

Controller

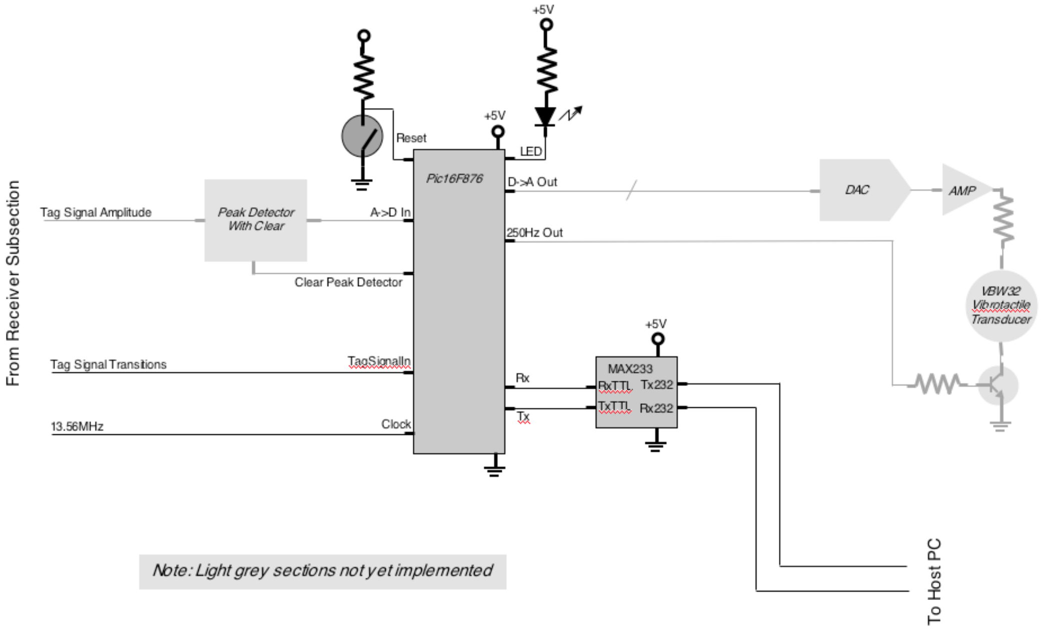

Transmitter & Antenna

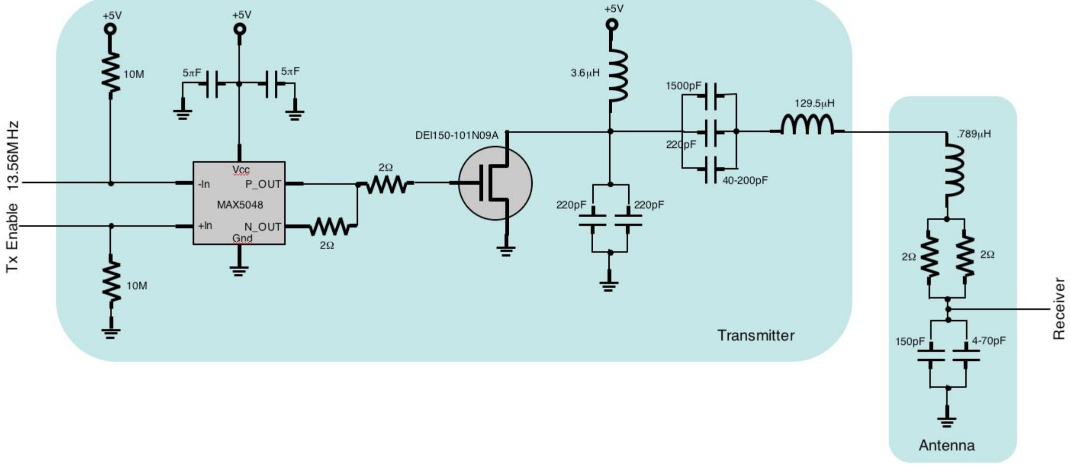

Appendix C. Microcontroller Code for Modified RFID Reader

(based on firmware provided with Microchip's 13.56MHz Evaluation kit.)

reader.asm

```
;Processor: PIC16f876 operating at 13.56 MHz
;                 Instruction time (Ti) = 295 nsec

            processor 16f876
            #include <P16F876.inc>
            __config h'3ff2' ;protection off,PWRT enabled,watchdog disabled,HS oscillator


;;INITIALIZATION PARAMETERS
;;++++++++++++++++++++++++++++++++++++++++++++++++++++++++
;Ports A,B,C
#define TRISA_INIT b'00000001'              ;clear all to outputs except b0 (analog signal input)
#define TRISB_INIT b'10000000'              ;clear all to outputs except b7 (digital signal input--interrupt source)
#define TRISC_INIT b'11000000'              ;clear all to outputs, set b7,b6 (serial rx,tx bit)

;Timers & Interrupts
#define OPTION_INIT b'00001000' ;PORTB pullups enabled (for outputs), INTEDG cleared,
                                            ;internal clock source for TMR0, T0SE cleared,
                                            ;prescaler assigned to WDT (means 1:1 prescaler for T0),
                                            ;prescaler code of 000
#define INTCON_INIT b'00001000'             ;disable G,P, enable RB,
                                            ;disable T0,INT interrupts
                                            ;clear T0F,INTF,RBIF flags
#define PIE1_INIT b'00000000'   ;disable PSP,AD,RC,TX,SSP,CCP1,TMR2,TMR1 interrupts
#define PIR1_INIT b'00000000'   ;clear PSP,AD,RC,TX,SSP,CCP1,TMR2,TMR1 flags
#define PIE2_INIT b'00000000'   ;disable EE,BCL, CCP2 interrupts
#define PE_IE INTCON,6                      ;for setting/clearing the PEIE (peripheral ie) bit
#define G_IE INTCON,7                       ;for setting/cearing the GIE (global ie) bit
#define RB_IE INTCON,RBIE       ;for setting/cearing the RBIE (port change ie) bit
#define T0_IF INTCON,T0IF       ;for setting/cearing the T0IF (timer0 if) bit

;Serial I/O
#define ClkFreq 13560000        ;external clock frequency = 13.56MHz
#define SPBRG_INIT d'87'        ;yields baud rate of 9600 for 13.56MHz clock
#define TXSTA_INIT b'00100100'              ;8-bit asynch transmission, hi speed, enable transmission
#define RCSTA_INIT b'10000000'              ;configure serial port pins
#define TX_EN TXSTA,5                       ;for setting/clearing the TXEN (transmit enable) bit
#define CR_EN RCSTA,4                       ;for setting/clearing the CREN (receive enable) bit
#define RC_IF PIR1,5                        ;for setting/clearing the RCIF (byte received) bit
#define TM_RT TXSTA,1           ;for checking if RS232 transmission has been completed

;A/D Input
#define ADCON0_INIT b'10000000'             ;conversion clock=Fosc/32, select channel 0, disable conversion
#define ADCON1_INIT b'00001110'             ;left-justify result, one analog input (RA0), Vdd&Vss refs
#define AD_IF PIR1,6                        ;A/D interrupt flag
#define AD_ON ADCON0,0                      ;A/D on bit
#define AD_GO ADCON0,2                      ;A/D conversion status bit
;;----------------------------------------------------


;;VARIABLE NAMES
;;++++++++++++++++++++++++++++++++++++++++++++++++++++++++
#define TAG_IN_A PORTA,0        ;tag signal (peak) input (analog) - is temporarily used as a digital output
                                            ;in a low state in order to drain the peak detector used to hold signal strength
                                            ;during A/D conversion
#define TAG_IN_D PORTB,7        ;tag signal input (digital)
#define DEBUG_OUT PORTA,5                   ;port pin for debugging
#define PWM_OUT PORTC,2                     ;PWM output (for a pseudo-analog output)
```

```
#define RS232_TX PORTC,6        ;serial port transmit pin
#define RS232_RX PORTC,7        ;serial port receive pin

#define BANK_0 PORTA                    ;hack to facilitate selecting bank 0
#define BANK_1 TRISA                    ;hack to facilitate selecting bank 1
#define BANK_2 EEDATA                   ;hack to facilitate selecting bank 2
#define BANK_3 EECON1                   ;hack to facilitate selecting bank 3

acctime =h'20' ;accumulated sync interval sum--also used as halfbit interval threshold
#define halfthr acctime ;halfbit interval threshold
prev_bit =h'21' ;the LSb stores the last rec'd bit--flip it by complementing f
bit_count = h'22' ;****
i = h'23' ;data byte iterator
j = h'24' ;data byte iterator
byte_to_translate = h'25'
offset = h'26' ;offset for table lookups
sum_bytes = h'27'
cs_difference = h'28'
fzero_test_result = h'29'       ;framing zeros are actually zero test result flag
cs_test_result = h'2a'          ;checksum test result flag
low_nibble_char = h'2b'
high_nibble_char = h'2c'
inner_count = h'2d'
outer_count = h'2e'
temp = h'2f'
adc_test_result = h'30'
signal_strength = h'31'


;;;!!!!!!!!!!!!!!!!!!!!!bit storage area--16 bytes of storage, indirectly addressed
;;;Note that s/w tests for MSb to detect end of area--be careful if move to different
;;;processor or relocate this storage area

received_tag_data_ptr = h'40'   ;32 bytes set aside for storing the received bits--actual number of bytes
                                ;in transmission is 4
#define NUM_TAG_DATA_BYTES h'04'    ;Number of (unprocessed) tag data bytes
                                    ;used in our tag data format
                                    ;(The tag actually sends 154 bits ~ 18 bytes)

;;Note that main loop uses bit tests to determine bit receive or runaway condition (to limit
;;processing time). Keep this in mind if received_tag_data_ptr storage area changed in the future.
;;
;;40h-60h is reserved for received bits--actual bit receiving area 40h-51h, rest is overrun area
;;52h-73h set aside for ASCII conversion of received bytes before RS232 transmission. Note that
;;52h-60h contains no useful information from the use during receive of demodulated bits. Also,
;; bits are not being received while the ASCII conversion and serial transmission are
;; taking place.
;; 'G' 1st character: "go"
;; Character 2-37: ASCII representation of received 18 bytes (until checksum used)
;; Character 38: '\n' newline

sendascii =h'52';begin of storage area for ASCII conversion of received bytes
xfercnt =d'14'  ;defines number of received bytes to convert to ASCII & transmit
;;-------------------------------------------------------


        org h'000' ;RESET vector location
        goto Start
        org h'004' ;interrupt service routine (ISR) vector location


;;INTERRUPT SERVICE ROUTINE
;;
;; Executes upon a PORTB<7:4> pin logic level transition (RB Port Change Interrupt)
;;
;; 1. BEWARE! To minimize interrupt response time, the w & status registers are NOT
;; archived.
;; 2. The ISR execution path is determined by the contents of the w register and
;; uses calculated goto's.
;; The value of w for the next ISR call is set at the end of the current ISR call
;; and is dependent on signal context (i.e. sync start, w/in sync, w/in data, etc.)
;; BEWARE! Must stay w/in 255 instructions for this to work!
;; 3. Sync field processed as follows:
;;    a. Ignore the first 4 transitions, they may be in response to tag power on reset.
```

```
;;    b. Accumulate the sum (time) of the next 8 intervals.
;;    c. Establish the width (time) of half a bit using the threshold width (time) for a
;;       full bit calculated as an average of the 8 interval measured above.
;;          (Repeating the previous bit, in a Manchester encoding, results in a sequence
;;          of two time periods--each with the width of half a bit--with an intermediary
;;          logic level transition.  Complementing the previous bit, on the other hand,
;;          results in a sequence of two time periods--each with the width of half a bit--
;;          with NO intermediary logic level transition.)
;;       The (time) width of half a bit is defined as 1.5x the average synch.
;;    d. Wait for a (time) width that is greater than the full bit width threshold. This is end of sync.
;;
;; BEWARE! The ISR presumes we are in Bank 0.
;;
;; In accordance with Manchester encoding, the sync field will be: 1 1 1 1 1 1 1 1 0
;;+++++++++++++++++++++++++++++++++++++++++++++++++++++++++
ISR:
            addwf PCL,f             ;4 calculated goto
                                    ;first sync edge is calculated goto here
            clrf TMR0 ;5
            movf PORTB,f ;6 must read PORTB before clearing RBIF
            bcf INTCON,RBIF ;7 just in case timer interrupt happened just at 1st edge
            bcf INTCON,T0IF ;8
            movlw (first_cycle - ISR - d'1') ;9 calculated goto offset for next ISR
            clrf prev_bit ;10 prev_bit @ end of sync = 0
            retfie      ;12
                        ;end of first cycle here. Note that first 4 transitions are ignored, because sync start is
                        ;corrupted by tag power on reset.

first_cycle
            clrf TMR0 ;5
            movf PORTB,f ;6 must read PORTB before clearing RBIF
            bcf INTCON,RBIF ;7
            movlw (second_cycle - ISR-d'1') ;8 calculated goto offset for next ISR

            ;Configure TAG_IN_A as an input pin to receive signal strength measure
            ;+
            banksel TRISA           ;10
            bsf TRISA,0             ;11
            banksel BANK_0          ;13  move back to the default bank, Bank 0
            ;-

            retfie      ;15
                        ;end of 2nd cycle here. Note that first 4 transitions are ignored, because sync start is
                        ;corrupted by tag power on reset.

second_cycle
            clrf TMR0 ;5
            movf PORTB,f ;6 must read PORTB before clearing RBIF
            bcf INTCON,RBIF ;7
            movlw received_tag_data_ptr ;8
            movwf FSR ;9 set up to store data bits
            movlw (third_cycle - ISR-d'1') ;10 calculated goto offset for next ISR

            ;start up A/D Converter - Need to wait at least 5 tag transmission edge changes
            ;for conversion to complete
            ;+
            bcf AD_IF ;11
            bsf AD_ON               ;12
            bsf AD_GO               ;13
            ;-

            retfie      ;15
                        ;end of 3rd cycle here. Note that first 4 transitions are ignored, because sync start is
                        ;corrupted by tag power on reset. The 3rd cycle is the 4th transition, so from here we measure
                        ;the longest interval in sync field.
third_cycle
            clrf TMR0 ;5
            movf PORTB,f ;6 must read PORTB before clearing RBIF
            bcf INTCON,RBIF ;7
            clrf acctime ;8 reset accumulated sync interval for average
            movlw (fourth_cycle - ISR-d'1') ;9 calculated goto offset for next ISR
            retfie      ;11
                        ;end of 4th cycle here. Start looking for longest sync interval here.
```

```
fourth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 first measured sync cycle, must be the largest
        movlw (fifth_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie      ;12
                ;end of 5th cycle here.

fifth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0
        movlw (sixth_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie      ;12
        ;end of 6th cycle here.

sixth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0
        movlw (seventh_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie ;12
        ;end of 7th cycle here.

seventh_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0
        movlw (eighth_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie      ;12
                ;end of 8th cycle here.

eighth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0
        movlw (nineth_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie      ;12
                ;end of 9th cycle here.

nineth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0

        ;Ensure A/D result is legitimate (finished), store it in signal_strength
        ;+
        btfss AD_IF             ;10 If A/D hasn't finished yet, there's been some mistake.
        incf adc_test_result,f  ;11
        movfw ADRESH            ;12
        movwf signal_strength           ;13
        ;-

        movlw (tenth_cycle - ISR-d'1') ;14 calculated goto offset for next ISR
        retfie  ;16
                ;end of 10th cycle here.

tenth_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
```

```
        addwf acctime,f ;9 acctime = acctime + TMR0

        ;Turn A/D Converter off
        ;+
        bcf AD_ON              ;10
        ;-

        movlw (eleventh_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie       ;13
                ;end of 11th cycle here. --this is last of sync cycles to be accumulated. Average the result
                ;and determine halfbit threshold in remaining sync cycles.
eleventh_cycle
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        addwf acctime,f ;9 acctime = acctime + TMR0
        movlw (twelfth_cycle - ISR-d'1') ;10 calculated goto offset for next ISR
        retfie      ;12
                ;end of 12th cycle here. Start averaging the sync interval accumulated time

twelfth_cycle
        movf PORTB,f ;5
        bcf INTCON,RBIF ;6
        rrf acctime,f ;7 acctime/2
        rrf acctime,f ;8 acctime/4
        rrf acctime,f ;9 avg interval = acctime/8
        movlw h'1f' ;10 clear 3 MSbs that may have been set by carry
        andwf acctime,f ;11
        movlw (cycle13 - ISR-d'1') ;12 calculated goto offset for next ISR
        retfie      ;14
                ;end of 13th cycle here. Calculate the halfbit threshold = 1.5(sync interval avg) Note that
                ;that the threshold value will be kept in acctime (=halfthr)

cycle13
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        rrf acctime,w ;8 half the sync interval avg
        addwf acctime,f ;9
        incf acctime,f ;10  acctime = halfthr = 1.5*(sync interval avg) + 1
        movlw (sync_end - h'100'-h'1'-ISR) ;11 calculated goto offset for next ISR
        bsf PCLATH,0 ;12 adjust for origin @ 100h
        retfie ;14

        org h'100'

;sync end wait. End of sync is distinguished by a fullbit interval. ( T > halfthr )
sync_end
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to detect end of sync field (halfthr - w)
        movlw (sync_end - h'100'-ISR-d'1') ;10 calculated goto offset for next ISR
        btfss STATUS,C ;12 Carry set for halfthr >= w
        movlw (bit1 - h'100'-ISR-h'1');12 If T > halfbit, end of sync detected. Proceed to data processing
        retfie ;14

;rec'd bit processing here --bit1 is 1st bit of 8 bit block
bit1
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit1 ;12

;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit2 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
```

```
        retfie ;17

halfabit1

        ;repeated bit (1 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half21-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie      ;17

;2nd half of bit interval processing
half21 ;2nd half, bit1
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit2-h'100'-ISR-h'1');8 calculated goto offset for next ISR
        retfie ;10
        ;rec'd bit processing here --bit2 is 2nd bit of 8 bit block

bit2
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit2 ;12

        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit3 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit2
        ;repeated bit (2 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half22-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17
        ;2nd half of bit interval processing

half22 ;2nd half, bit2
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit3-h'100'-ISR-h'1');8 calculated goto offset for next ISR
        retfie ;10
        ;rec'd bit processing here --bit3 is 3rd bit of 8 bit block

bit3
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit3 ;12
        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit4 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit3
        ;repeated bit (3 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half23-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

;2nd half of bit interval processing
```

```
half23 ;2nd half, bit3
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit4-h'100'-ISR-h'1');8 calculated goto offset for next ISR
        retfie ;10

;rec'd bit processing here --bit4 is 4th bit of 8 bit block
bit4
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit4 ;12

        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit5 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit4
        ;repeated bit (4 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half24-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

;2nd half of bit interval processing
half24 ;2nd half, bit4
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit5-h'100'-ISR-h'1');8 calculated goto offset for next ISR
        retfie ;10

;rec'd bit processing here --bit5 is 5th bit of 8 bit block
bit5
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit5 ;12

        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit6 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit5
        ;repeated bit (5 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half25-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

;2nd half of bit interval processing
half25 ;2nd half, bit5
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit6-h'100'-ISR-h'1') ;8 calculated goto offset for next ISR
        retfie ;10

;rec'd bit processing here --bit6 is 6th bit of 8 bit block
bit6
        movf TMR0,w ;5
```

```
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit6 ;12

;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit7 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit6
        ;repeated bit (6 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half26-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

;2nd half of bit interval processing
half26 ;2nd half, bit6
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit7-h'100'-ISR-h'1') ;8 calculated goto offset for next ISR
        retfie ;10

;rec'd bit processing here --bit7 is 7th bit of 8 bit block
bit7
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit7 ;12
        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
        movlw (bit8 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

halfabit7
        ;repeated bit (7 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half27-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17

;2nd half of bit interval processing
half27 ;2nd half, bit7
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit8-h'100'-ISR-h'1') ;8 calculated goto offset for next ISR
        retfie ;10

;rec'd bit processing here --bit8 is 8th bit of 8 bit block
bit8
        movf TMR0,w ;5
        clrf TMR0 ;6
        movf PORTB,f ;7
        bcf INTCON,RBIF ;8
        subwf halfthr,w ;9 Test interval to determine bit. C = 1 for repeated bit
        btfsc STATUS,C ;11
        goto halfabit8 ;12

        ;fullbit processing here
        comf prev_bit,f ;12 Complement prev_bit for fullbit measurement
        rrf prev_bit,w ;13
        rlf INDF,f ;14 shift in the new bit
```

```
        movlw (bit1 - h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        incf FSR,f ;16
        retfie ;18

halfabit8
        ;repeated bit (8 of 8)
        rrf prev_bit,w ;13
        rlf INDF,f ;14
        movlw (half28-h'100'-ISR-h'1') ;15 calculated goto offset for next ISR
        retfie ;17
        ;2nd half of bit interval processing

half28 ;2nd half, bit8
        clrf TMR0 ;5
        movf PORTB,f ;6
        bcf INTCON,RBIF ;7
        movlw (bit1-h'100'-ISR-h'1') ;8 calculated goto offset for next ISR
        incf FSR,f ;9 advance to next byte in received_tag_data_ptr storage array
        retfie ;11
;;--------------------------------------------------------


        org h'200'


;;INITIALIZATION
;;+++++++++++++++++++++++++++++++++++++++++++++++++++++++
Start:

        call Init_PortA
        call Init_PortB
        call Init_PortC
        call Init_Timer0
        call Init_ISR_Branching
        call Init_AD
        call Init_Serial_Port
        call Init_Interrupts


        ;;BEWARE! From here onwards, when listening for tag data & syncs, we stay in Bank 0
        ;;(to ensure that we are in Bank 0 for ISR)
        banksel BANK_0


SeekSyncInit:   ;initialization for sync field search
                ;executed on power-up or after data recovery
                ;has finished/failed

        clrwdt  ;clear watchdog timer

        ;Configure TAG_IN_A as an output pin of low state in order to drain leaky peak detector
        ;that holds signal strength measurement
        banksel PORTA
        bcf TAG_IN_A    ;set TAG_IN_A's corresponding output latch to low
        banksel TRISA
        bcf TRISA,0     ;configure TAG_IN_A as an output
        banksel BANK_0  ;move back to the default bank, Bank 0

        ;clear test result flags
        clrf cs_test_result
        clrf fzero_test_result
        clrf adc_test_result

        ;clear the signal strength estimate
        clrf signal_strength

        ;;clear the bit storage field
        ;;++++++++++
        movlw d'19'
        movwf bit_count ;initialize bit_count

        movlw received_tag_data_ptr         ;initialize tag data pointer
        movwf FSR
```

```
clrbits:
        clrf INDF               ;erase tag data byte
        incf FSR,f              ;increment pointer
        decfsz bit_count,f
        goto clrbits
        ;----------

        movlw received_tag_data_ptr             ;initialize tag data pointer
        movwf FSR

        ;read PORTB before clearing INTCON to be sure RBIF=0
        movf PORTB,w
        clrf INTCON
        clrf TMR0

        ;; Reset offset & PCLATH to 0
        movlw d'0'              ;w=0
        clrf PCLATH

        ;; Enable interrupt
        bsf RB_IE ;enable portB change interrupt
        bsf G_IE   ;global interrupts now enabled
;;--------------------------------------------------------


;;MAIN PROGRAM
;;
;;  BEWARE! In the main loop, w represents the PCL offset for calculated gotos in the ISR.
;;  (don't use it for other stuff unwittingly)
;;
;; The main loop monitors the T0IF flag to detect successfully received word (subject to
;; checksum test). Tag word processing is ISR driven. ***->A calculated goto method is used for
;; position context in tag word for speed.
;;
;; If the main loop detects a timer overflow, the w register is cleared to return processing
;; to the search for a first sync edge.
;;
;; Also, expect received_tag_data_ptr area to be @ 40h-52h **** while receiving data.
;; This must be checked bitwise (because w can't be used in the main loop).
;;+++++++++++++++++++++++++++++++++++++++++++++++++++++++++


SeekSync:
        bcf RB_IE
        movlw d'0' ;calculated goto offset for 1st sync edge processing

        clrf PCLATH

        clrf FSR ;FSR = 0 to indicate not gathering bits

        bsf RB_IE
        bcf T0_IF

Main:
        clrwdt      ;clear watchdog timer

        btfsc FSR,6
        goto ReceivingTagData ;receiving data, monitor progress

        btfsc INTCON,T0IF
        goto SeekSync ;if TMR0 overflows w/o receiving bits, seeksync

        goto Main

ReceivingTagData:
        clrwdt      ;clear watch dog timer

        btfsc T0_IF
        goto SeekSyncInit       ;if timer overflows,***WHAT HAPPENS HERE?****

        btfsc FSR,2             ;If bit 2 set, FSR > 43h and the first 4 data bytes of tag transmission
                                ;have been stored.  Stop receiving, start processing.
        goto ProcessTagData
        goto ReceivingTagData
```

```
ProcessTagData:

        ;;First, disable interrupts (thoroughly!)
        ;;++++++++++
        clrf INTCON
clrgie:
        bcf INTCON,GIE
        btfsc INTCON,GIE ;make sure it's clear before proceeding
        goto clrgie
        movf PORTB,f
        clrf INTCON ;disable all interrupts while processing received data
        ;;----------

SetExtraZeroTimer:
        ;;Next, set a timer so we can keep track of when its ok to start looking for the next tag
        ;;transmission. (Tags bit-streams are 154 bits long and we only make use of the first 36 bits.
        ;;This means that 118 bits of junk must be sent before we can begin looking for the
        ;;next tag transmission.  Since the tags have a minimum data rate of 58KHz, each bit
        ;;takes a maximum time of 17.24usec and the total required wait time for the
        ;;118 bits is ~.00204 seconds.)
        ;;++++++++++
        ;assign prescaler to Timer0, set prescaler to 32
        movlw b'00000100'
        banksel OPTION_REG
        movwf OPTION_REG

        ;preload timer with (256-217) to get a wait time of Ti*prescaler*217 = .002048sec
        movlw d'39'
        banksel TMR0
        movwf TMR0

        ;clear T0 interrupt flag
        banksel INTCON
        bcf T0_IF

        ;move to the default bank, Bank 0
        banksel BANK_0
        ;;----------

DrainPeakDetector:
        ;Configure TAG_IN_A as an output pin of low state in order to drain leaky peak detector
        ;that holds signal strength measurement
        banksel PORTA
        bcf TAG_IN_A            ;set TAG_IN_A's corresponding output latch to low
        banksel TRISA
        bcf TRISA,0             ;configure TAG_IN_A as an output
        banksel BANK_0          ;move back to the default bank, Bank 0

RemoveFramingZeros:
        ;;Remove the framing '0' bits from the received databytes & checksum
        ;;by bit shifting the data array left until all framing 0s are shifted out
        ;;++++++++++++++++++++
        ;initialize iterators & a bit counter
        movlw NUM_TAG_DATA_BYTES - 1
        movwf i                 ;external iterator
                                ;(i changes once per pass through all bytes-to-be-rotated)
        movlw NUM_TAG_DATA_BYTES
        movwf j                 ;internal iterator (j changes once for each rotation of a byte)
        clrf bit_count          ;bit counter (holds the sum of all the framing zeros.
                                ;bit_count *should* hold 0 after the framing zeros are removed.
        bcf STATUS,C            ;clear carry bit (its rotated in during a rotation)

rm_fzero_pass:
        ;set FSR to address of final tag data byte
        movlw received_tag_data_ptr + NUM_TAG_DATA_BYTES
        movwf FSR

        ;At this point, C holds the framing bit that was rotated out on the last rotation of the
        ;previous pass.  If the framing bit is *not* a zero, the data is corrupt.
        ;Increment bit_count to keep a count of non-zero framing bits.
        btfsc STATUS,C
        incf bit_count,f
```

```
        bcf STATUS,C            ;clear carry bit
                                ;(it's value gets rotated into final data byte during first rotate command.)

rotate_byte:

        ;decrement j. if j==0 goto done_rm_fzero_pass, else goto in_rb
        decfsz j,f
        goto in_rb
        goto done_rm_fzero_pass

in_rb:
        ;FSR--
        decf FSR,f

        ;left-rotate the current byte
        rlf INDF,f

        goto rotate_byte

done_rm_fzero_pass:

        ;decrement i. if i==0 goto done_rm_fzeros, else set j=i+1 and goto rm_fzero_pass
        decfsz i,f
        goto in_drmfzp
        goto done_rm_fzeros

in_drmfzp:
        movfw i
        movwf j
        incf j,f
        goto rm_fzero_pass

done_rm_fzeros:
        ;At this point, C holds the final framing bit. Add it to the bit count.
        btfsc STATUS,C
        incf bit_count,f

        ;;Test whether or not the framing zeros are in fact zero.
        incf bit_count,f ;bit_count should now be 1
        decfsz bit_count,f
        goto fzero_error
        goto done_checking_fzeros

fzero_error:
        incf fzero_test_result,f

done_checking_fzeros:
        ;;--------------------

CalcChecksum:
        ;;Calculate checksum. Sum the [first...penultimate] tag data bytes, compare
        ;;sum to the final byte. Calculation only accounts for the (low) 8 bits.
        ;;+++++++++++++++++++
        movlw received_tag_data_ptr     ;move to the beginning of the tag data
        movwf FSR

        movlw ((NUM_TAG_DATA_BYTES - 1) - 1)
        movwf i                 ;i holds the number of bytes left to sum

        clrf sum_bytes          ;clear sum_bytes to hold the running sum

add_byte:
        movfw INDF
        addwf sum_bytes,f

        incf FSR,f

        decfsz i,f
        goto add_byte

test_checksum:
        bcf STATUS,C            ;clear carry bit

        movfw INDF              ;move transmitted checksum to w
```

```
        subwf sum_bytes,f     ;subtract transmitted checksum from calculated checksum,
        movfw sum_bytes               ;place result in cs_difference
        movwf cs_difference

        btfss STATUS, C               ;if C is clear, (sumbytes - transmitted checksum) < 0 => checksums differ
        goto checksums_differ
        incf cs_difference,f   ;(preload for decfsz test)
        decfsz cs_difference,f ;if (sumbytes - transmitted checksum) != 0 checksums differ
        goto checksums_differ
        goto checksums_match
checksums_differ:
        incf cs_test_result,f
checksums_match:
        ;;--------------------

ExamineResults:
        ;; If the tag data passes framing zero & checksum tests, report a successful tag read - otherwise don't.
        ;;++++++++++++++++++++
        ;;Does data pass the adc, framing zero, checksum tests?
        incf adc_test_result,f        ;(preload for decfsz tests)
        incf fzero_test_result,f      ;(preload for decfsz tests)
        incf cs_test_result,f         ;(preload for decfsz tests)
        decfsz adc_test_result,f      ;if fzero_test_result!=0, tag data is bad
        goto DiscardResults
        decfsz fzero_test_result,f    ;if fzero_test_result!=0, tag data is bad
        goto DiscardResults
        decfsz cs_test_result,f       ;if cs_test_result!=0, tag data is bad
        goto DiscardResults
        goto ReportResults

DiscardResults:
        ;goto ExtraZerosTimeOut

ReportResults:

        bsf DEBUG_OUT

        ;Send tag data via RS232
        movlw d'13'
        call RS232_Send_Byte
        movlw d'10'
        call RS232_Send_Byte

        movlw a'a'
        call RS232_Send_Byte
        movfw adc_test_result
        movwf byte_to_translate
        call RS232_Send_ASCII

        movlw a'z'
        call RS232_Send_Byte
        movfw fzero_test_result
        movwf byte_to_translate
        call RS232_Send_ASCII

        movlw a'c'
        call RS232_Send_Byte
        movfw cs_test_result
        movwf byte_to_translate
        call RS232_Send_ASCII

        movlw d'13'
        call RS232_Send_Byte
        movlw d'10'
        call RS232_Send_Byte

        movlw h'40'
        movwf FSR
        movfw INDF
        movwf byte_to_translate
        call RS232_Send_ASCII

        movlw h'41'
```

```
		movwf FSR
		movfw INDF
		movwf byte_to_translate
		call RS232_Send_ASCII

		movlw h'42'
		movwf FSR
		movfw INDF
		movwf byte_to_translate
		call RS232_Send_ASCII

		movlw a' '
		call RS232_Send_Byte

		movfw signal_strength
		movwf byte_to_translate
		call RS232_Send_ASCII

		movlw d'10'
		call RS232_Send_Byte
		movlw d'13'
		call RS232_Send_Byte
		;;--------------------

ExtraZerosTimeOut:
		;;Here we wait for the tag to finish sending (unused) extra zeros, then reinitialize Timer0 for
		;;receiving the next tag transmission.
		;;++++++++++
extra_zeros_wait_loop:
		btfss T0_IF		;will be set by timer overflow
		goto extra_zeros_wait_loop

		;reinitialize Timer0 for next tag read
		movlw OPTION_INIT
		banksel OPTION_REG
		movwf OPTION_REG

		;move back to the default bank, Bank 0
		banksel BANK_0
		;;----------

DoneTagRead:
		;;All done.  Start looking for the next tag transmission.
		bcf DEBUG_OUT
		goto SeekSyncInit

SpinCycle:
		goto SpinCycle
;;------------------------------------------------------


;;INITIALIZATION FUNCTIONS
;;++++++++++++++++++++++++++++++++++++++++++++++++++++++
;;Initialize Port A
Init_PortA:
		banksel TRISA
		movlw TRISA_INIT
		movwf TRISA

		banksel PORTA
		movlw d'0'
		movwf PORTA

		return

;;Initialize Port B
Init_PortB:
		banksel TRISB
		movlw TRISB_INIT
		movwf TRISB

		banksel PORTB
		movlw d'0'
		movwf PORTB
```

```
		return

;;Initialize Port C
Init_PortC:
		banksel TRISC
		movlw TRISC_INIT
		movwf TRISC

		banksel PORTC
		movlw d'0'
		movwf PORTC

		return

;;Init Timer0
Init_Timer0:
		banksel OPTION_REG
		movlw OPTION_INIT
		movwf OPTION_REG

		return

Init_ISR_Branching:
		banksel PCLATH
		movlw HIGH ISR
		movwf PCLATH

		return

;;Initialize A/D converter
Init_AD:
		banksel ADCON0
		movlw ADCON0_INIT
		movwf ADCON0

		banksel ADCON1
		movlw ADCON1_INIT
		movwf ADCON1

		return

;;Initialize Serial Port
Init_Serial_Port:
		banksel SPBRG		;reset & initialize baud rate generator
		movlw SPBRG_INIT	;set up for 9600 baud rate
		movwf SPBRG

		banksel RCSTA		;set up for serial reception
		movlw RCSTA_INIT
		movwf RCSTA

		banksel TXSTA		;set up for serial transmission
		movlw TXSTA_INIT
		movwf TXSTA

		banksel BANK_0
		bsf CR_EN		;enable serial port reception

		banksel BANK_1
		bsf TX_EN		;enable serial port transmission

		return

;;Initialize Interrupts
Init_Interrupts:

		banksel INTCON
		movlw INTCON_INIT
		movwf INTCON

		banksel PIE1
		movlw PIE1_INIT
		movwf PIE1
```

```
        banksel PIR1
        movlw PIR1_INIT
        movwf PIR1

        banksel PIE2
        movlw PIE2_INIT
        movwf PIE2

        return
;;-------------------------------------------------------


;RS232 FUNCTIONS
;;+++++++++++++++++++++++++++++++++++++++++++++++++++++++
;Receive Byte - receives contents into w register
RS232_Receive_Byte:
        bcf RC_IF ;clear byte received flag
        movf RCREG, w       ;load received data into w
        bsf G_IE   ;ensure interrupts are enabled
        bsf PE_IE  ;ensure interrupts are enabled
        return

;Send Byte - transmits contents of w register
RS232_Send_Byte:
        banksel BANK_0
        movwf TXREG                 ;load data to send. transmission starts automatically

        banksel BANK_1              ;TM_RT is in bank 1
in_sending:                         ;wait for transmission to finish
        btfss TM_RT                 ;wait until TMRT is set
        goto in_sending

        banksel BANK_0
        return

;RS232 Send G - transmits an ascii 'G'
RS232_Send_G:
        movlw a'G'
        call RS232_Send_Byte
        return

;RS232 Send ASCII - transmits the ascii equivalent of a byte (two characters)
RS232_Send_ASCII:

        ;;Isolate low nibble, calculate ASCII equivalent
        ;;++++++++++++++++++++
        movfw byte_to_translate     ;load byte_to_translate into w

        andlw h'0f'         ;isolate the LSN
        addlw d'1'  ;increment (table entries start at an offset of 1)
        movwf offset

        bcf STATUS,C        ;clear C - its used later to detect page crossings

        movlw LOW Hex2ASCII         ;get low 8 bits of table address

        addwf offset,f              ;add them into the offset

        movlw HIGH Hex2ASCII        ;get high 5 bits of table address

        btfsc STATUS,C      ;page crossed?
        addlw d'1'          ;yes, then increment high address

        movwf PCLATH        ;load high address in PCLATH latch

        movfw offset        ;load computed offset into w

        call Hex2ASCII      ;table lookup, ends with ASCII char in w
        movwf low_nibble_char
        ;;--------------------

        ;;Isolate high nibble, calculate ASCII equivalent
        ;;++++++++++++++++++++
```

```
            rrf byte_to_translate,f ;rotate MSN to LSN positions
            rrf byte_to_translate,f
            rrf byte_to_translate,f
            rrf byte_to_translate,f
            movfw byte_to_translate         ;load (modified) byte_to_translate into w

            andlw h'0f'             ;isolate the LSN

            addlw d'1'  ;increment (table entries start at an offset of 1)
            movwf offset

            bcf STATUS,C            ;clear C - its used later to detect page crossings

            movlw LOW Hex2ASCII             ;get low 8 bits of table address

            addwf offset,f                  ;add them into the offset

            movlw HIGH Hex2ASCII            ;get high 5 bits of table address

            btfsc STATUS,C          ;page crossed?
            addlw d'1'              ;yes, then increment high address


            movwf PCLATH            ;load high address in PCLATH latch

            movfw offset            ;load computed offset into w

            call Hex2ASCII          ;table lookup, ends with ASCII char in w
            movwf high_nibble_char
            ;;--------------------

            ;;Send the two ascii characters, high nibble first
            movfw high_nibble_char
            call RS232_Send_Byte
            movfw low_nibble_char
            call RS232_Send_Byte

            return
;;--------------------------------------------------------


            org h'3ff'


;Hexadecimal to ASCII conversion table
;;+++++++++++++++++++++++++++++++++++++++++++++++++++++++
Hex2ASCII
            addwf PCL,f
            retlw "0" ;ascii 0
            retlw "1" ;ascii 1
            retlw "2" ;ascii 2
            retlw "3" ;ascii 3
            retlw "4" ;ascii 4
            retlw "5" ;ascii 5
            retlw "6" ;ascii 6
            retlw "7" ;ascii 7
            retlw "8" ;ascii 8
            retlw "9" ;ascii 9
            retlw "A" ;ascii A
            retlw "B" ;ascii B
            retlw "C" ;ascii C
            retlw "D" ;ascii D
            retlw "E" ;ascii E
            retlw "F" ;ascii F
;;--------------------------------------------------------


            end
```

Appendix D. Lessons from Implementing an RFID Reader

In addition to practical lessons bearing upon the design of interactive behaviors for a hand-held locator, I learned practical lessons useful for realizing such a device. The present Appendix briefly summarizes what I discovered to be the main technical challenges of building an RFID reader subsystem by subsystem, then presents an invaluable list of "tips and tricks".

Transmitter: Efficiency proved to be the main technical challenge of the transmitter subsection. In practice, it can be difficult to build and tune an RF power amplifier with efficiency over 70% (though with practice and the right components, efficiencies of over 90% can be obtained) (Sokal, 2000). Efficiency is especially important in the context of a high-powered, wireless hand-held device, since battery life is finite and lost power becomes heat close to the hand.

Power: Efficiency and power density are the main technical challenges of the power sub-system, for the same reasons discussed in reference to the transmitter subsection.

Receiver: The main technical challenge of the receiver subsection is extracting relatively weak tag transmissions from a strong carrier signal. In the context of a hand-held device, an RFID receiver shares close quarters with a powerful transmitter, and so precautions must be taken to isolate it through careful shielding, grounding and filtering.

Control: The interrupt service routine for handling tag data must execute and return very quickly in order to keep up with a tag's data rate of 70kbits/sec. I used the transmitted 13.56MHz carrier as a clock signal and a PIC16F876 microprocessor to interpret incoming tag data, and found that the fastest possible interrupts were just barely fast enough to handle tag signal transitions. Aside from this concern, implementing the control subsection is relatively straightforward for anyone familiar with firmware.

Numerous traps lie in store for the would-be implementer without a strong background in RF electronics. The following tips and tricks are included so that these traps may be circumnavigated by future investigations:

- The subsystems of an RFID reader are highly interdependent, so don't obsess over the performance of any one subsystem. Get the whole system working, then tweak subsystems with an eye towards the performance of the complete system.
- The signal generated by the transmitter is powerful; it can impact measurement if precautions are not taken. Position lab equipment and circuitry carefully, and design test points within the circuit that minimize lengths of unshielded wire.
- Even at the intermediate frequency of 13.56MHz, standard approaches of linear electronics come under strain. Operational amplifiers must be selected with care. Realization of "ideal diodes" from operational amplifiers and discrete diodes is impractical.
- Since RFID tag data is decoded on the basis of transition timing, the shape of the incoming tag data waveform must not be significantly distorted before it is thresholded to reproduce the tag transponder's digital waveform. In the interest of minimizing distortion, filter stages should pass not only the frequency corresponding to the tag's data rate, but also several *harmonics* of the this data rate (this is and important consideration for filtering a digital signal that is typically not mentioned in classical texts on filter design).
- Axial approximations for the magnetic field of a coil antenna suggest a maximum field strength at the coil's center in the plane of the coil, however the magnetic field may in practice be stronger near the coil's periphery. As a result, approximations of tag-reader distance based on axial approximations of magnetic field strength may begin to break down when tag and reader are extremely close to one another.
- Given the aim of maximizing an antenna's magnetic field strength (to extend read range), an antenna based on a series resonator is preferable to an antenna based on a parallel resonator. (Commercial systems tend to employ parallel

resonators primarily because they can be matched more easily to 50Ω transmission lines of arbitrary length).

- Of the various power amplifier topologies that exist, the "Class E" topology is particularly well suited to the transmitter subsection of an RFID reader. This topology is efficient at intermediate and high frequencies, can be tuned visually, and requires only one (switching-mode) power transistor.
- It is tempting to consider use of the logarithmic AM demodulators and Received Signal Strength Indicators (RSSI) developed for mobile phones to extended the dynamic range of an RFID receiver subsection. One must keep in mind, however, that RFID transmissions rest upon an extremely strong carrier signal, while cellular transmissions do not.
- Traditional electronic prototyping and evaluation techniques can be slow, error-prone and noise-prone. Computer-aided design, simulation and layout tools can facilitate the process tremendously. (I relied extensively on one CAD tool – HEPA PLUS from Design Automation[40] – to good effect for re-designing the transmitter subsection, and would recommend availing of CAD tools for other subsections as well).
- Sourcing parts for the transmitter subsection and antenna can be particularly difficult, owing to strict simultaneous requirements for power, frequency and size. Some of my key finds were:
    - The MAX5048 High-Speed High-Current MOSFET Driver with adjustable rise and fall times, from Maxim Integrated Products. It proved useful as a small and efficient bridge between a high-speed low-power crystal oscillator and a high-speed, high-power transistor.
    - The DE150-101N09A RF Power MOSFET, from Directed Energy. This high-frequency power transistor's low input capacitance, low drain-source resistance, high drain-source breakdown voltage and high drain-source current capability make it ideal for 13.56MHz RFID applications.
    - CWS Bytemark's T68-2 Carbonyl 'E' Iron Powder Toroidal Inductor Cores. These cores proved useful for constructing compact inductors with a low AC resistance at 13.56MHz. (Good inductors at

intermediate and high frequencies are made rather than bought, and making them well requires a way to measure their properties. Appendix E describes the apparatus and technique I used to measure the resistance and inductance of inductors at 13.56MHz.)

- A.C.L Mica Capacitors, from Europa Components & Equipment. The low resistance and high tolerances of these devices made them useful for constructing resonators, and for measuring inductance experimentally. (Discussion of techniques for tuning resonators appears in Appendix E)
- RTO-Series Thick Film Non-Inductive Power Resistors, from Vishay. Resistors employed to tune the Q of an antenna resonator at 13.56MHz must be selected for high power, low inductance and accurate resistance values. These resistors were well suited to the task.
- Adhesive Copper EMI Sheilding tape, from Chomerics. Adhesive tape made from copper is useful for constructing High-Q antennas; it has low electrical resistance, and, unlike the copper tubing typically used in antenna construction, it can be placed, removed, bent and cut rapidly and without specialized tools.

---

[40] See http://www.techexpo.com/firms/desauto.html

Appendix E. Apparatus and Procedure for Inductor Measurement and Resonator Tuning.

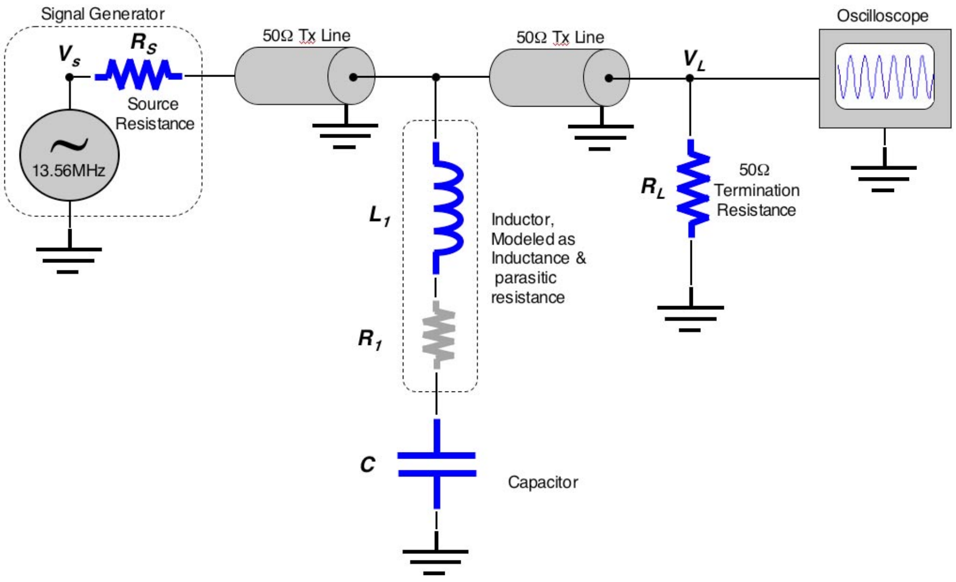

Measuring an Inductor’s Inductance:

Adjust the frequency of the signal generator, and watch the oscilloscope. Record the frequency corresponding to minimum signal amplitude. L_1 can be determined from

$$f_0 = \frac{1}{2\pi\sqrt{LC}}$$

Measuring an Inductor’s Resistance:

Adjust the frequency of the signal generator, and watch the oscilloscope for the minimum signal amplitude f_0. Record the minimum signal amplitude V_L the amplitude of the source signal V_S. R_1 can be determined from

$$R_1 = \frac{V_L R_L}{V_S - 2V_L}.$$

Tuning a Resonator:

Set the frequency of the signal generator to the desired resonant frequency. Adjust the value of *C* to achieve minimum signal amplitude on the oscilloscope display.

Appendix F. Legality of a Hand-Held RFID locator with Range Exceeding a Meter.

FCC regulations in the US (Part 15.225) stipulate that power transmitted at 13.56MHz should not exceed 10mV/m at a distance of 30m from the transmitter (Lee and Sorrells, 2001). In Europe, ETSI regulation EN300330 stipulates that transmitted power at a frequency of 13.56MHz should not exceed 42dBuA/m at a distance of 10m from the transmitter (Finkenzeller, 2004). The calculations of this appendix suggest that an RFID reader with a 0.2m diameter antenna coil can successfully provide power to credit-card sized tags at a distance >1 meter while remaining within legal limits. (Calculations presume a 7.2V, 2.1A power supply, and a DC-RF power conversion efficiency of 70%). The MATLAB script used to run the calculations and its associated output file appear below.

SCRIPT

```
%Power, Range & Legality Calcs for Hand-Held RFID Tag Reader
% Goal: Design an RFID tag reader for hand-held use that is
% * small
% * light
% * long-lasting
% * entirely wireless (power & signal)
% * has a long range
% * reports signal strength

%This file holds the power and range calculations
%Variables

% Vcell - voltage of an individual cell used in the battery powering the reader
% Rcell - resistance of an individual cell
% num_cells -number of cells in battery
% Vbatt - voltage of the battery
% Rbatt - resistance of battery
% ITbatt - current/time rating (eg 1700mAH)
% RunTime - time for which system much work before a battery recharge is necessary
% Vreg - voltage of the regulated source directly powering the circuit
% RegulationEfficiency - efficiency of the regulator for the power amplifier
% Ireglim1 - maximum current from the regulated source given voltage drop from Vbatt over Rbatt
% Ireglim2 maximum current from the regulated supply Imax as limited by required RunTime
% Ireg - maximum current the regulated source can supply at Vreg for Tbatt
%       given voltage drop from Vbatt over Rbatt
% Irxc  - upper limit on current required by receiver section
% Icontrolc - upper limit on current required by control section
% Iconvc - current available to DC->RF converter circuit
% Pconvc - DC power available to DC->RF converter circuit
% ConversionEfficiency - Efficiency of DC->RF conversion
% Prf - RF power delivered to the conversion circuit's load
% Prfrms - rms value associated with Prf
% Rload - load (antenna) resistance
% Irfrms - rms value of current passing through converter's load
% N - number of turns in antenna coil
% Bo - minimum B field necessary to power tag
% miu0 - permeability of free space
```

```
% a - radius of coil
% r - maximum read range

diary results.txt
%Power is really the ultimate constraint on this system, so we start there.

%Define the battery
Vcell = 1.2; %v
Rcell = .0065+.01; %ohms  %.0065 max internal resistance of a sanyo sub-c cell (web);
%.01 a guess at solder tags r...
num_cells = 6;
Vbatt = num_cells * Vcell %v
Rbatt = num_cells * Rcell %ohms
ITbatt = 1.2 %Ah

%Choose a run time
RunTime = .4 %hours; normal for 7.2V R/C car batteries

%Set a voltage for the regulated voltage supply (less than Vbatt, big enough to
%drive logic)
Vreg = 5 %v
RegulationEfficiency=.7; %possible with switched cap regulator...

%Set a maximum current for the (ideal) regulated supply based on two constraints.

%First, the current shouldn't cause the voltage of the regulated supply to "dip"
Ireglim1 = (Vbatt - Vreg)/Rbatt

%Second, the current should be small enough so that the battery can supply it for the duration of the system runtime
Ireglim2 = ITbatt/RunTime
Ireg = RegulationEfficiency*min(Ireglim1,Ireglim2)

%Now in order to work in the Tx, Rx and control sections with some independence, we have to

%do some "current budgeting".  Ideally, we'd like to send all the power to the Tx
%section, but the other sections need power too.  Lets set upper limits for the current
%requirements of the other sections, then give what is left to the Tx section.
Irxc = .150; %A max current required by receiver section
Icontrolc = .050; %A - max current required by control section
Iconvc = Ireg - Irxc - Icontrolc %A - current available to DC->RF converter circuit

%Calculate the power available to the antenna, assuming a value for the DC->RF conversion
%efficiency.
ConversionEfficiency = .7 %possible with class e
Pconvc = Vreg*Iconvc; %W DC power available to DC->RF converter.
Prf = Pconvc * ConversionEfficiency %W Power actually converted to RF.
Prfrms = Prf %The power can be interpreted as RMS power. ****

%The method for choosing the component parameters of a "class E" DC->RF generator
%fixes the load resistance:
Rload = 3 %ohms

%What is the current flowing through this resistance?
%Re-arranging and combining Vrms=Irms*R & Prms=R*Irms^2 we have:
Irfrms = sqrt((Prfrms)/Rload) %A

%Now, what sort of read range can this current generate?
%Here we assume that
%  * a tag requires .0449 * 10^-6 wbm^-2 (Microchip 13.56MHz RFID System Guide p79)
%  * the resistance is the load (no matching loss)

N = 5; %antenna coil turns
tag_area = .0046224; %m^2
res_freq = 13.56*1000000; %Hz
N_tag = 6; %turns
Q_tag = 40;
CoilTurnOnVoltage = 4; %Vpp
angle_tag = 0; %radians
Bo = ((4/sqrt(2))/(2*pi*res_freq*N_tag*tag_area*Q_tag*cos(angle_tag))) %wbm^-2
%Bo = .0449*10^-6 %wbm^-2
a = .1 %m - radius of coil
miu0 = 4*pi*(10^-7); %H/m; pemeability of free space

%Solve for r iteratively; done when Ireqrms exceeds Irfrms
```

```
prevIreqrms=0;
Ireqrms=0;
r=0;
r_step = .001;
while(Ireqrms<=Irfrms)
  r=r+r_step;
  prevIreqrms = Ireqrms;
  Ireqrms = (2*Bo*(a^2+r^2)^(3/2))/(N*miu0*a^2);
end;

%We've gone one past; move one back by calculation
r=r-r_step
Ireqrms = prevIreqrms
Ireqpeak = Ireqrms*sqrt(2)

%Now, is it legal?

%FCC regulation 15.255 stipulates a maximum power level at fundemental frequency
%of 10mV/m @ 30m from transmitter
Haxialrms = (Ireqrms * N * a^2)/(2* sqrt((a^2+30^2)^3)); %A/m
%Conversion equations from http://www.radioing.com/eengineer/convert.html
%dBmV = 20 log [Signal (mV)/1mV]
%dBmA = 20 log [Signal (mA)/1mA]
%dBmA/M = dBmV/M - 51.5

%(Where the constant 51.5 is a conversion of the characteristic
%impedance of free space (120p) into decibels: 20Log10[120p] = 51.5)
Haxialrms_dBmAperM = 20*log10(1000*Haxialrms); %dBmA/m
Haxialrms_dBuAperM = 20*log10(1000000*Haxialrms); %dBuA/m
Haxialrms_dBmVperM =Haxialrms_dBmAperM+51.5; %dBmV/m
Haxialrms_mVperM = 10^(Haxialrms_dBmVperM/20); %mV/m

%(these calcs were verified using "RF Units" from Texas Instruments RFID)
mVperMat30M = Haxialrms_mVperM
if (Haxialrms_mVperM<10)
   'legal in USA!'
else
   'not legal in USA.'
end

%ISM regulations stipulate a maximum power level at fundemental frequency
%of 42dBuA/m @ 10m from transmitter
%http://www.rfid-handbook.de/rfid/frequencies.html
Haxialrms = (Ireqrms * N * a^2)/(2* sqrt((a^2+10^2)^3)); %A/m
%Conversion equations from http://www.radioing.com/eengineer/convert.html
%dBmV = 20 log [Signal (mV)/1mV]
%dBmA = 20 log [Signal (mA)/1mA]
%dBmA/M = dBmV/M - 51.5
%(Where the constant 51.5 is a conversion of the characteristic
%impedance of free space (120p) into decibels: 20Log10[120p] = 51.5)
Haxialrms_dBmAperM = 20*log10(1000*Haxialrms); %dBmA/m
Haxialrms_dBuAperM = 20*log10(1000000*Haxialrms); %dBuA/m
dBuAperMat10M = Haxialrms_dBuAperM

if (Haxialrms_dBuAperM<42)
   'legal in EU!'
else
   'not legal in EU.'
end

diary off
```

OUTPUT

```
Vbatt =
   7.2000

Rbatt =
   0.0990

ITbatt =
```

```
   1.2000

RunTime =
   0.4000

Vreg =
    5

Ireglim1 =
  22.2222

Ireglim2 =
   3.0000

Ireg =
   2.1000

Iconvc =
   1.9000

ConversionEfficiency =
   0.7000

Prf =
   6.6500

Prfrms =
   6.6500

Rload =
    3

Irfrms =
    1.4888

Bo =
    2.9924e-008

a =
    0.1000

r =
    1.1560

Ireqrms =
    1.4880

Ireqpeak =
    2.1044

mVperMat30M =
    0.5178

ans =
    legal in USA!

dBuAperMat10M =
    31.4096

ans =
    legal in EU!
```

Appendix G. Form-Giving: Sketches & Doodles

The sketches and doodles below were created in order to open up a space where the form of a hand-held locator might be considered in reference to activities similar to locating through indirect assistance. Activities such as illuminating, truffle hunting, magnifying, aiming and listening to heartbeats suggested a variety of directions form might take.

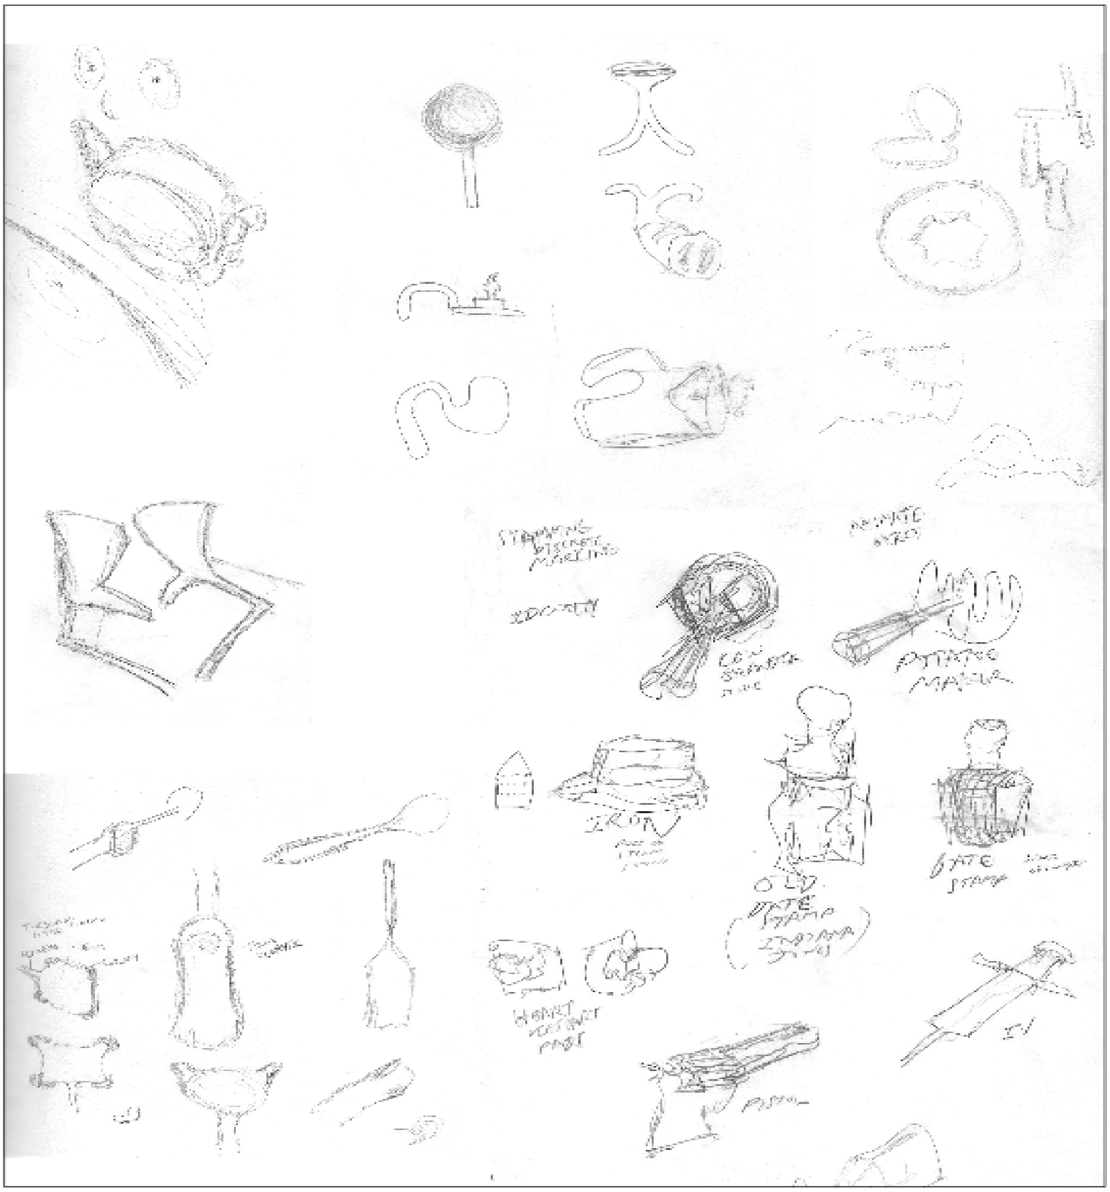

Appendix H. Verifying the Linearity of the Vibrotactile Transducer

Before using the vibrotactile transducer in experiments, it seemed prudent to verify that it converted electrical energy to vibration linearly. If it didn't, *designed* and *delivered* vibrotactile gradients would differ.

To verify linearity, I fastened three small, lightweight accelerometers (ADXL202's from Analog Devices) orthogonally to the transducer, and monitored their reported acceleration values while varying the amplitude of the software-generated 250Hz drive signal. The drive signal was varied over a range comparable to the range used in the simulation, and the transducer was mechanically "loaded" with a pressing thumb. Since the drive signal was a sinusoid, amplitude of acceleration could be used in lieu of amplitude of displacement. (Displacement is the double integral of acceleration).

After recording a number of <software-set vibrotactile amplitude, actual magnitude of acceleration> data points, I graphed the results. The graph to the left shows that the magnitude of acceleration is, in fact, linear with input voltage. The graph to the right shows that the component of acceleration corresponding to transducer movement towards and away from the thumb is a) linear and b) greater than the components of acceleration corresponding to lateral movement against the thumb pad. (The psychophysical studies I encountered on digit sensitivity to vibrotactile stimuli were all based on axial vibration towards and away from the digit pad).

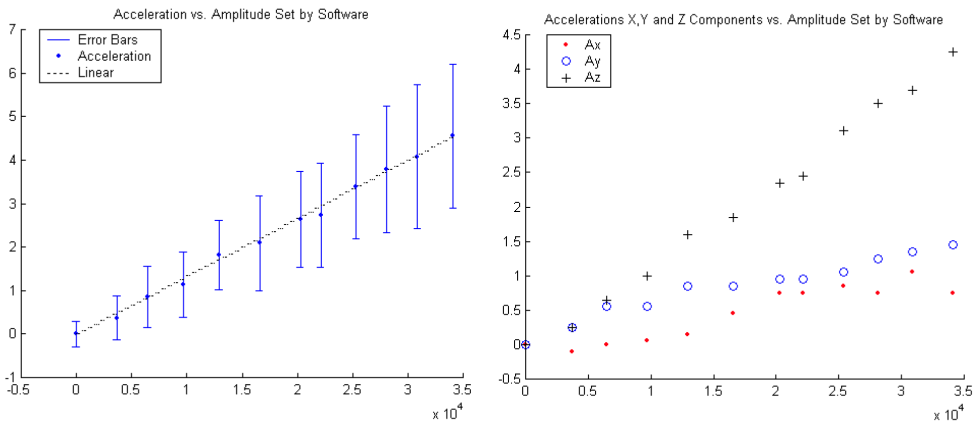

Left: Magnitude of acceleration as a function of signal amplitude set in software. The relationship is approximately linear. Right: Magnitude of acceleration's components as a function of software-set signal amplitude. The component corresponding to movement towards and away from the thumb is the strongest component, and is approximately linear.

Matlab script used to generate the graphs for linearity test results

```
%m-file to graph the linearity test results for the
%vibrotactile transducer
%V: output voltage of the power amplifier
%Aud, Ass, Aoi: components of transducer's measured acceleration
%S_a,b vibrotactile amplitude set in software


%SET UP DATA

V = [0 .463 .946 1.19 1.805 2.415 3.2 3.65 4.11]
S_a = [0 4148 8295 10139 15208 20277 27190 30416 34102];
V_err = [V(1:3).*.2 V(4).*.15 V(5).*.12 V(6:end).*.05]
V_slope = (V(end)-V(1))/(S_a(end)-S_a(1))
V_lin = V_slope.*S_a

S_b = [0 3687 6452 9678 12904 16590 20277 22120 25346 28111 30877 34102]
Aud = [.25 .5 .9 1.25 1.85 2.1 2.6 2.7 3.35 3.75 3.95 4.5]
Ass = [.25 .15 .25 .3 .4 .7 1 1 1.1 1 1.3 1.0]
Aoi = [.25 .5 .8 .8 1.1 1.1 1.2 1.2 1.3 1.5 1.6 1.7]
%Get rid of base-level noise
Aud = Aud - .25
Ass = Ass - .25
Aoi = Aoi - .25
%Merge acceleration’s components
Aall = sqrt(Aud.^2+Ass.^2+Aoi.^2)
%Calculate error
Aud_err = [.1 .2 .3 .3 .3 .4 .4 .5 .5 .75 .75 .75]
Ass_err = [.1 .1 .1 .15 .2 .3 .3 .3 .3 .3 .4 .4]
Aoi_err = [.1 .2 .3 .3 .3 .4 .4 .4 .4 .4 .5 .5]
Aall_err = [Aud_err+Ass_err+Aoi_err]
%Calculate linear approximations
Aud_slope = (Aud(end)-Aud(1))/(S_b(end)-S_b(1))
Aud_lin = Aud_slope.*S_b + Aud(1)
Ass_slope = (Ass(end)-Ass(1))/(S_b(end)-S_b(1))
Ass_lin = Ass_slope.*S_b + Ass(1)
Aoi_slope = (Aoi(end)-Aoi(1))/(S_b(end)-S_b(1))
Aoi_lin = Aoi_slope.*S_b + Aoi(1)
Aall_slope = (Aall(end)-Aall(1))/(S_b(end)-S_b(1))
Aall_lin = Aall_slope.*S_b + Aall(1)


%GRAPH RESULTS

hold off
figure('Name','Voltage','Color','white')
hold on
errorbar(S_a,V,V_err,V_err,'b.')
plot(S_a,V_lin,'k:')
axis([0 35000 0 5])
title('Voltage into Vibrator vs. Amplitude Set by Software');
legend('Error Bars','Voltage','Linear')

hold off
figure('Name','Acceleration Components','Color','white')
hold on
%errorbar(S_b,Ass,Ass_err,Ass_err,'r.')
%plot(S_b,Ass_lin,'r:')
plot(S_b,Ass,'r.')
%errorbar(S_b,Aud,Aud_err,Aud_err,'b.')
%plot(S_b,Aud_lin,'b:')
plot(S_b,Aoi,'bo')
%errorbar(S_b,Aoi,Aoi_err,Aoi_err,'k.')
%plot(S_b,Aoi_lin,'k:')
plot(S_b,Aud,'k+')
%axis([])
title('Accelerations X,Y and Z Components vs. Amplitude Set by Software');
%legend('Error Bars x','Actual Ax','Linear Ax')
%legend('Error Bars y','Actual Ay','Linear Ay')
%legend('Error Bars z','Actual Az','Linear Az')
legend('Ax','Ay','Az')
```

```
hold off
%subplot(2,2,3)
figure('Name','Acceleration','Color','white')
hold on
errorbar(S_b,Aall,Aall_err,Aall_err,'b.')
plot(S_b,Aall_lin,'k:')
%axis([])
title('Acceleration vs. Amplitude Set by Software');
legend('Error Bars','Acceleration','Linear');
```

Appendix I. Latency Measurement

The latency of vibrotactile response was determined through sonic means. I substituted a speaker for the vibrotactile element, and positioned the locator near to the edge of a target's active region (the region where vibrotactile feedback would occured). When the locator was tapped vigorously, it got "bumped" into the active region. The sound of the tap as well as the onset of the 250Hz response were recorded, and the waveform examined in DigiDesign's ProTools (a sound application). Through inspection, it was possible to discern the delay between the tap and the onset of the 250Hz response.

This was repeated four times, with latencies of .07, .09, .088 and .088 seconds. The rounded average latency, .09 seconds, was in keeping with the report rate of the RFID tags the system was simulating (10/second), so no additional delays needed to be introduced in software. (Ideally there would be no delay, however the aim was to create a faithful simulation of a feasible implementation, and this feasible implementation did entail delay).

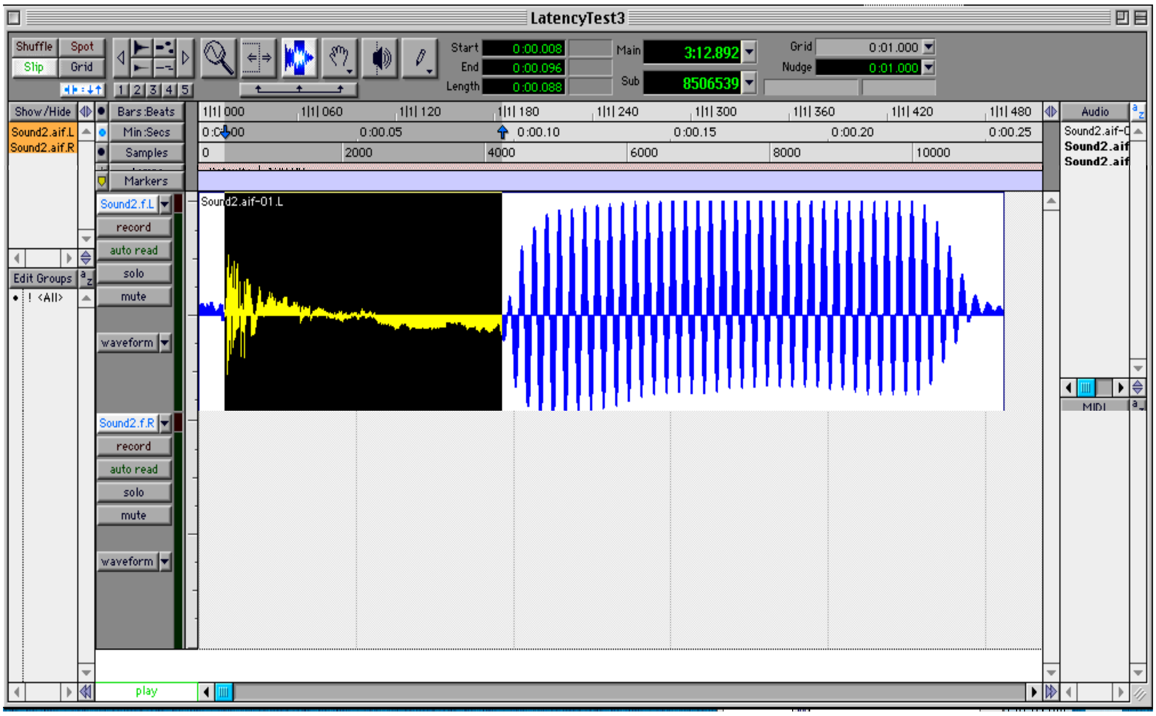

Start 0:00.008
End 0:00.096
Length 0:00.088
Main 3:12.892
Sub 8506539
Grid 0:01.000
Nudge 0:01.000

Waveform from the “tap test” used to determine system latency. The tap appears near the beginning of the waveform, the onset of the 250Hz vibration appears approximately .09 seconds later.

Appendix J. Script for 1st Experiment

Are you familiar with a dowsing rod? A metal detector? A Geiger counter?

What I'm doing is trying to invent a hand-held device similar to these things, but for finding items labeled with these tags **<show RFID tags>**. These are a bit like the tags used on CDs to prevent shoplifting.

Thats the Big Project; and right now I'm just focusing on one tiny piece: designing the feedback the "locator" provides to the person using it.

There are lots of ways the locator could communicate "closer" or "farther"; it could use a light that gets brighter as it gets closer; it could use a sound that gets louder. It could become hotter like in the children's game. I'm concentrating on vibration – “vibrotactile feedback” -- felt by the thumb of the hand holding the locator.

Here. Try moving this slider widget with this mouse. What do you feel?
<verify that participant experiences a changing amplitude of vibration>

So thats the basic idea: have this sensation you feel become more intense when your locating hand gets close, and less intense as you get farther away.

Here is a simulation in 2D space with a mouse **<show simulation>**. The mouse pointer is for locating a target. The target is that grey circle there. Try moving the mouse pointer slowly toward the circle. What do you feel?

(As mouse pointer moves toward the circle, it enters range of the function mapping amplitude of vibration to distance. As the mouse pointer moves closer, vibrotactile feedback grows according to the selected linear mapping function. As the mouse pointer enters the target, the feedback begins to pulse on and off)

<verify that participant experiences these things>

Right, so you can feel that intensity grows as it gets closer? Well, this can happen in different ways.

<Change the mapping function from linear to threshold>

Now try. Do you feel how its different?

<Verify that the participant feels a difference>

You can imagine that *how* the intensity changes with distance from the target makes a difference in trying to find the target. Here; give this a try.

<Make the target invisible and move it to a random point on the screen>

Now the target is invisible. Can you find it?

<When the participant finds it, change the mapping function and target position again>

Can you find it now?

<after target has been found the second time>

Was it easier to find the target the first time or the second time?

Ok, so that's what this experiment is about: trying to see how the way intensity changes with distance influences the experience of finding a target.

Now we are going to move from this 2D virtual space with the mouse into a 3D space with a prototype of the locator **<show the 3D space and locator. Place a**

virtual target in the 3d space>. Can you find the target now? **<Show how to hold the locator device>**.

This experiment is a bunch of time trials. You start with the locator here **<indicate starting position>**. After moving to the start position, you find the virtual target in this space here **<indicate the 3d space>**. When you have found it – when the vibration begins to pulse – press the button on the locator. Be careful to press the button firmly. Pressing the button ends the time trial. When you are ready to start the next time trial, move the locator back to the starting position.

Now, for a few practice runs. **<allow a few practice runs>**

Do you have any questions?
Are you ready to do the trials?

There will be 96 trials and two intermissions, 1/3 and 2/3 of the way through (cookies will be provided). It will probably take about 20 minutes. If you think any of the trials should be discounted (because you were interrupted, had to scratch your knee, etc) write the trial number down on this paper here. If you get lost, just follow the instructions in the prompt window.

You probably noticed that the vibrotactile device makes sound as well. Since this experiment focuses only on vibrotactile feedback, please wear these ear muffs while doing the trials.

When you are done, please turn the paper over and fill out the questionnaire.

I'll be watching and video taping (if this is ok). Thanks for your help & happy targeting!

Appendix K. Questionnaire to be Completed After the First Experiment

What is your
name?
contact email address / phone number?
height?
age?
gender?
occupation?
handedness?

Are you familiar with using a computer & mouse?

How did you find the target? What strategies did you try?

During this experiment you were presented with several patterns of vibrotactile feedback. How many kinds do you remember? Please describe them. How were they alike/different?

Did you prefer one pattern of vibrotactile feedback over the others? Why?

Was this experience of trying to find a location through vibrotactile feedback like anything you've done before? If so, what was it like?

In what situations or occupations could you envision this capability being useful?

What would you tell the next person doing this experiment in order to better prepare them for it?

Appendix L. Script for 2nd Experiment

Your participation in the last experiment helped to determine how the intensity of a vibrotactile stimulus can change with distance in order to facilitate rapid targeting of an invisible point in 3d space. **<Show the 4 functions that were tested>**. From the results of this experiment, the "Perceived Gaussian" mapping function was chosen as the "fastest" function.

This second experiment takes things one step further. It is about how the *range* of the "best" function affects the ease of targeting.

You can see **<demonstrate with the 2D mouse simulation>** that when the range is very small, the feedback isn't very useful. The same is also true when the range is very big **<demonstrate with the 2D mouse simulation>.** If the range is very big, it becomes difficult to follow the gradient toward the target. If the range is not too big and not too small, *that's* where the feedback is most useful for targeting. The purpose of this experiment is to find out about this middle ground and to relate it to target size.

The format of this experiment is similar to the last one: a bunch of time trials. You start with the locator here **<indicate starting position>**. The circle on the prompt window will turn yellow when you are inside the start position **<indicate the circle on the prompt window>**. After moving to the start position, you try to find the virtual target in this space here **<indicate the 3d space>**. When you have found it – when the vibration begins to pulse – press the button on the locator. Be careful to press the button firmly. Pressing the button ends the time trial. When you are ready to start the next time trial, move the locator back to the starting position.

Now, for a few practice runs. **<allow a few practice runs>**

Do you have any questions?
Are you ready to do the trials?

There will be 96 trials and two intermissions, 1/3 and 2/3 of the way through (cookies will be provided). It will probably take about 20 minutes. If you think any of the trials should be discounted (because you were interrupted, had to scratch your knee, etc) write the trial number down on this paper here. If you get lost, just follow the instructions in the prompt window.

You probably noticed that the vibrotactile device makes sound as well. Since this experiment focuses only on vibrotactile feedback, please wear these ear muffs while doing the trials.

When you are done, please turn the paper over and fill out the questionnaire.

I'll be watching and video taping (if this is ok).

Thanks for your help & happy targeting!

Appendix M. Questionnaire to be Completed After the Second Experiment

(will be accompanied by first questionnaire for new participants)

Did the range of the feedback change during the experiment? If so:
How many different feedback ranges do you recall being presented with?
Did you find it easier to locate the target when the range was larger or when it was smaller?

Did the target change size? If so:
How many different size targets do you recall being presented with?
Did you find it easier to locate the target when the range was larger or when it was smaller?

In this experiment, you experienced how vibrotactile feedback from a hand-held device could be used to locate a *virtual* target. If you could locate *physical* items in shelf or room sized spaces with a similar (wireless) device, what would you use it for?

Where would the capability of finding through vibrotactile feedback be useful?

Based on your experience of this study, how would you recommend that this investigation proceed? (What would you see as the "next step"?)
Thanks for your help.

Appendix N. Additional Questionnaire Responses

First Questionnaire

What would you tell the next person doing this [first] experiment in order to better prepare them for it?

Participants tended to offer these four pieces of advice:

- Start with gross, fast gestures and successively refine them to find the target.
- Remember that the target could just as easily be near the center of the search volume as near its borders; don't cling to the walls of the volume.
- Don't worry about breaking the elastic string; it isn't going to break.
- Don't get frustrated or disappointed if you can't find a particular target. Just move on.

Second Questionnaire

Where would the capability of finding through vibrotactile feedback be useful?

This question was repeated from first experiment, in hopes that participants may have considered potential applications for a hand-held locator. The application domains suggested were, for the most part, the same as previously suggested (e.g. darkrooms, libraries, chemical labs, stockrooms, warehouses, supermarkets). One participant noted that forgetfulness and misplacement are often issues faced by the elderly, and entrepreneurially hinted that the average age in countries such as the United States is on the rise.

In this [second] experiment, you experienced how vibrotactile feedback from a hand-held device could be used to locate a virtual target. If you could locate physical items in shelf or room sized spaces with a similar (wireless) device, what would you use it for?

This question was asked with the aim of eliciting applications that participants could speak about as experts, applications they could appreciate in their own lives. Aside from one participant who emphatically wanted support for finding his own keys, participants tended not address this question from their own personal perspective.

Appendix O. Declaration of Informed Consent for Participation in Time-Trial Experiments

I, the undersigned, hereby declare that I am willing to take part in a research project that is part of a Masters in Computer Science course experiment at the University of Limerick.

This study is entitled: Targeting through Vibrotactile Feedback. It examines vibrotactile feedback, felt through the thumb, as an aid for targeting sought items in shelf-sized spaces using a hand-held locator. The study involves first receiving instructions on the use of the experimental apparatus, then conducting a number of time trials. The duration of the time trials is approximately one hour. I understand that there will be two rest breaks during the trials, and that cookies will be provided.

I declare that I have been fully briefed on the nature of this study and my role in it and have been given the opportunity to ask questions before agreeing to participate. I understand that the purpose of this experiment is to evaluate the utility of a stimulus to assist targeting; the experiment is not an evaluation of my ability, knowledge or intelligence.

I understand that after performing the tasks I will be asked to answer several questions in relation to usability, realism, and suitability for performing these types of tasks, but that no personal, private, or confidential information will be required of me.

I fully understand that there is no obligation on me to participate in this study and that I am free to withdraw my participation at any time without having to explain or give a reason. I am also entitled to full confidentiality in terms of the details of my participation and my personal details. I understand that some or all of the data (verbal and behavioral) may be used (quoted) in the report on the evaluation for illustrative purposes, but that I shall not be identifiable from this data, either in the body of the report or in the appendices.

I also understand that my participation in this study may be recorded by video or audio means, as well as by notes taken by observers. I am entitled to copies of all records made during the session if so I wish to have them.

I acknowledge the fact that deception and concealment are inappropriate to, and not required in, this study, and that no attempt will be made to elicit information or actions from me using these means.

____

____

Signature of Participant   Date

Appendix P. Relative Frequency of Mapping Function Orderings, Data Split by Mapping-Range

The charts below depict the relative frequencies of (pair-wise) mapping function orderings across subjects for two values of mapping-range (18cm and 32cm) and two levels of statistical significance (75% and 95% confidence intervals). On these charts, scores for the "less than" condition indicate participants for whom the first mapping function correlates with statistically shorter find-times than the second mapping function (a statistically significant difference). Scores for the "greater than" condition indicate participants for whom the first mapping function correlates with longer find-times than the second mapping function (a statistically significant difference). Scores for the "equals" condition indicate participants for whom the first and second mapping function are statistically equivalent. All comparisons were made using a pair-wise Wilcoxon signed rank tests.

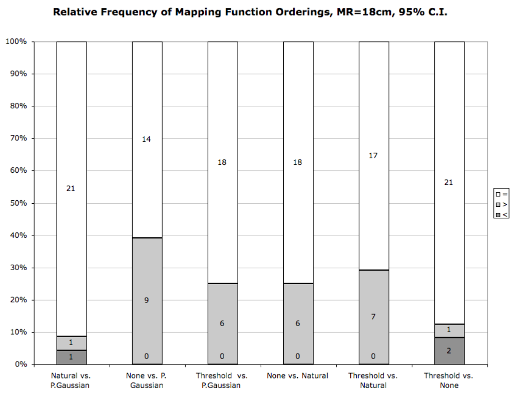
Relative Frequency of Mapping Function Orderings, MR=18cm, 95% C.I.

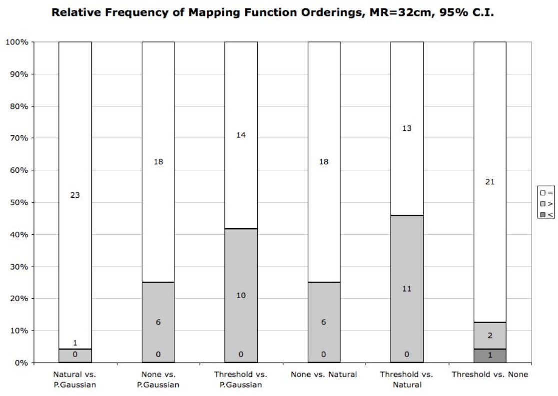
Relative Frequency of Mapping Function Orderings, MR=32cm, 95% C.I.

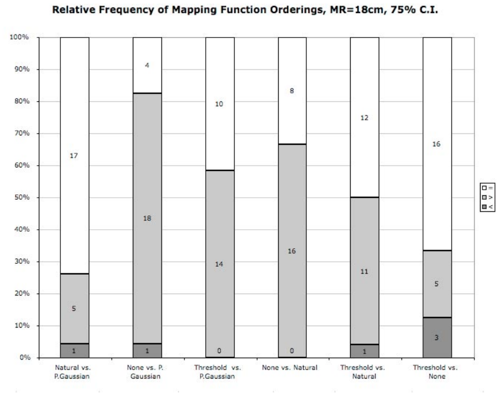
Relative Frequency of Mapping Function Orderings, MR=18cm, 75% C.I.

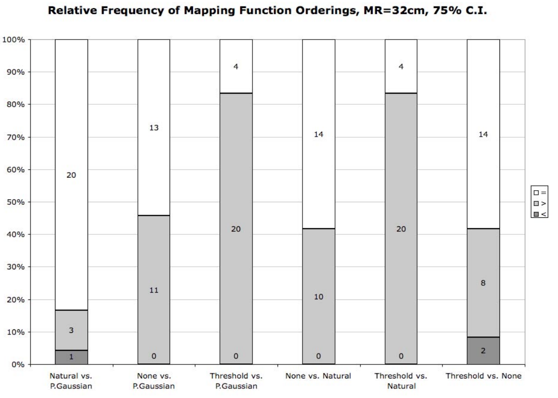
Relative Frequency of Mapping Function Orderings, MR=32cm, 75% C.I.

At a confidence interval of 95%, the white regions—the regions corresponding to no statistical difference in median find-time between the two mapping functions—constitute the majority for all comparisons between mapping functions (for both values of mapping-range). At a confidence interval of 75%, the same white regions are in the minority for most comparisons, but are still relatively large. Lowering the confidence interval decreases the frequency of ties for comparisons of both mapping range values, but does not significantly alter the remaining balance between greater-thans and less-thans; this suggests that there are trends in the data pointing towards faster and slower mapping functions, but that the statistical tests are too stringent for the trends to be conclusive with the data collected. Perhaps with greater sample sizes and more tightly controlled experiments, the trends would present themselves more conclusively.