Towards the Distributed Visualization of Usage History

Paul Michael Yarin

BSE Mechanical Engineering, University of Pennsylvania, 1995
MS Mechanical Engineering, Stanford University, 1997

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences
at the Massachusetts Institute of Technology

September 1999

Author — **Program in Media Arts and Sciences**
August 6, 1999

Certified by — **Hiroshi Ishii**
Professor of Media Arts and Sciences
Program in Media Arts and Sciences

Accepted by — **Stephen A. Benton**
Chair, Departmental Committee for Graduate Students
Program in Media Arts and Sciences

Towards the Distributed Visualization of Usage History

Paul Michael Yarin

**Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning, on August 6th, 1999,
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences
at the Massachusetts Institute of Technology**

Abstract

Because they require focused visual attention and explicit user control, current computer interfaces have either limited computer augmentation to highly specialized environments or have necessitated major changes to work practice. "Situated usage history displays," devices that indicate the former use of physical objects and surfaces enable convenient access to information without disruption of existing environments or behavior. Prototype systems for monitoring the usage of handheld objects, storage containers, medication bottles and other devices are presented; underlying technologies for sensing and display are also described. The design space of situated usage history displays is characterized through discussion of key opportunities and constraints.

**Thesis Advisor:
Hiroshi Ishii
Professor of Media Arts and Sciences
Massachusetts Institute of Technology**

Towards the Distributed Visualization of Usage History

Thesis Committee

Thesis Reader		Hiroshi Ishii Associate Professor of Media Arts and Sciences
Thesis Reader		Mitchel Resnick Associate Professor of Media Arts and Sciences
Thesis Reader		Scott Snibbe Member, Research Staff, Interval Research Corporation

Your notes here

Acknowledgments

Hiroshi Ishii
For providing an engine, premium fuel, and a heavy foot on the pedal

Scott Brave, Andy Dahley, Matt Gorbet, Craig Wisneski, Victor Su, Phil Frei, Jay Lee, Seungho Choo, and James Patten
For being exciting, creative, and fearless colleagues

Josh Smith, Tom White, Rich Fletcher, Golan Levin, and Jon Ferguson
For sharing time and talent as collaborators

Hiroshi Ishii, Mitch Resnick, and Scott Snibbe
For their efforts to review and improve this document

Clay Harmony and James Hsiao
For their diligent project execution

David, Naoko, and Peter Yarin
For the gift of a happy and loving family

Joanna Berzowska
For her insight and charm

Brygg Ullmer and John Underkoffler
For being omnivorous in intellect, if not in diet

John Maeda, Neil Gershenfeld, Marvin Minsky, Henry Jenkins, Larry Leifer, Ed Carryer, Dennis Boyle, Hale Rowland, Tommy Cowan, and Bob Goody
For teaching me well

Scamper
For teaching me to be a dog, well

Table of Contents

1	INTRODUCTION	8
	1.1 Motivation	8
	1.2 Research Topic	9
	1.3 Thesis Overview	10
2	BACKGROUND AND RELATED WORK	11
	2.1 Computer-augmented environments	11
	2.2 Situated information spaces	12
	2.3 History-enriched digital objects	13
	2.4 Tangible interfaces	14
	2.5 Artificial intelligence, perception, and emotion	15
3	DESIGN APPROACH	17
	3.1 Objectives	17
	3.2 Design Parameters	17
4	GESTURAL INTERACTION DEVICES	20
	4.1 Introduction	20
	4.2 FishFace	20
	4.3 ShakePad	25
	4.4 Discussion	26
5	AUGMENTED CONTAINER STORAGE SYSTEM	29
	5.1 Background	29
	5.2 Implementation	30
	5.3 Interaction techniques	31
	5.4 Discussion	34

6	OBJECT IDENTIFICATION AND TRACKING SURFACES	37
	6.1 Introduction	37
	6.2 SensePad	39
	6.3 IDPad	45
	6.4 Discussion	49
7	EVALUATION AND FUTURE WORK	51
	7.1 FishFace and ShakePad	51
	7.2 TouchCounters	52
	7.3 SensePad and IDPad	52
8	CONCLUSION	54
	8.1 Summary	54
	8.2 Personal vision	54
A	APPENDIX: ADDITIONAL INTERACTIVE DEVICES	56
	A.1 Tactile Pad	56
	A.2 Network Meter	59
B	APPENDIX: PROJECT INFLUENCE MAP	61
C	APPENDIX: DESIGN PARAMETERS	62
	REFERENCES	63

1 Introduction

1.1 Motivation

The basic operations of work are exchanges between people, information, and physical objects (Cooper98). Although computers have greatly facilitated the ease and efficacy of interpersonal communication, similar benefits have yet to be widely conferred upon interactions with physical objects.

One reason is that most computer interfaces are poorly suited for mediating interactions with physical objects. The dominant paradigm for computer interaction is the "graphical user interface" or GUI, in which users monitor a detailed display of text and graphics while typing on keyboards and manipulating pointing devices. While well suited to applications in which the computer is the focus of attention, this type of interaction is difficult to carry out in conjunction with inherently physical tasks—the sorting of items on a table, the location of physical objects or containers, or the use of physical tools.

Under these circumstances, users stand to benefit most from "lightweight" interfaces, those with which data can be usefully recorded or retrieved with minimal cognitive and physical effort. To enable lightweight interactions with information, interfaces should be

- physically present and visible at the location of use
- easy to comprehend through use of visual patterns
- responsive to simple physical actions
- capable of conveying relevant information

The successful design of these systems demands more than simply extrapolating the design of the GUI to larger and smaller form factors. Lightweight interfaces depend upon the careful integration of sensing technology, display, and interaction techniques in a physical form suitable to the environment of use. Simplicity and efficiency, rather than raw computational power determine real-world utility.

1.2 Research Topic

To address the above goals, I have chosen to develop various computational devices that demonstrate a possible approach. These devices aim to enhance collaboration by displaying visual representations of locally recorded physical data. Unlike the cumbersome, generic computers of today, these devices are envisioned as integral and unobtrusive features of living and working environments, acquiring and displaying information silently and automatically. Because these devices portray the past histories of objects and environments, their means of interaction may be termed “the distributed visualization of usage history.”

This phrasing requires some clarification. *Distributed* refers not to a method of data analysis but rather to the simultaneous exchange of data at multiple physical locations. *Visualization* refers to the visual representation of information in a manner that facilitates understanding. *Usage* refers broadly to any physical activity performed repeatedly to achieve a useful function. *History* signifies data that has been accumulated over time, whether or not the temporal dimension is explicitly represented to the user.

The concept of usage frequency as an aid to collaboration has roots in both online media, in its use of relevance queries and popularity rankings, as well as in physical indications of wear: walking trails worn by repeated passage, fingerprints on the popular pages of repair manuals, and date stamps in the backs of popular library books (Hill92). This work concerns electronically synthesizing such phenomena in response to physical actions.

The arguments in favor of measuring usage history include its flexibility, simplicity, and robustness. The measurement of usage frequency applies equally to characterizing the histories of either people or objects. The many available sensing technologies make detection of a wide variety of physical events feasible at a low cost, and with a high level of reliability. (Note that this technique may, but need not, take advantage of remotely stored, computed, or acquired data; networking is not a prerequisite for measurements of usage history.) As measurement of history can be readily automated, these systems can acquire data without the intervention of a human operator.

Distributed visualization describes both the presentation of information from multiple, individual displays, and the techniques that allow data to be compared through concurrent viewing of multiple displays in parallel. Techniques for information visualization are well established in the computer graphics community (Card99), but are rarely applied to augmenting of physical objects or spaces with multiple displays.

Displays designed to integrate with physical structures can be positioned within the visual field of the user. This greatly reduces the cost of attending to the interface, as setup and tear-down times become essentially nonexistent. In addition, the use of multiple displays in parallel allows identification of trends in relative use across objects and spaces; this is especially well suited to objects contained in furniture with regular geometries like shelves, drawers, or pegboards.

1.3 Thesis Overview

This document presents prototypes and conceptual models of situated usage history displays. Related work in computer-augmented environments, information visualization, and digital histories are described. In the next section, the design space for situated usage history displays is outlined through several key parameters. The development of the several prototype systems are then presented. The next chapter evaluates these systems and describes recommendations for future work. Finally, a brief summary is presented with a description of a personal vision for this work.

The first Appendix described two additional prototypes that fall somewhat outside the scope of this work. The second charts the conceptual connections between projects; the third lists the design dimensions in greater detail.

2 Background and Related Work

This work draws upon a variety of influences and research findings in computer-augmented environments, context-aware computing devices, interaction histories for software objects, information visualization, and tangible interfaces. Related projects not discussed here include augmented reality (Feiner93), distributed sensing devices (Poor99), and computational learning tools (Resnick98).

2.1 Computer-augmented environments

A wide variety of projects claim to address the integration of computing technologies with physical spaces. Most, however, assume that people exchange the same *types* of information in shared physical spaces as they do when using desktop computers—messages, meeting notes, indications of attendance, etc. As a result, most systems tend to incorporate scaled versions of desktop-style graphical user interfaces—pointing devices, icons, multimodal, and high-resolution displays.

Ubiquitous computing

The Ubiquitous Computing concept (Weiser91) anticipated the widespread proliferation of networked computing devices throughout interior environments. Wall-sized projected whiteboards and tablet-sized computers were demonstrated, in addition to a handheld device called the ParcTab. The later companion concept of "calm technology" (Weiser97) emphasized the importance of subtle and peripheral indicators of human activity.

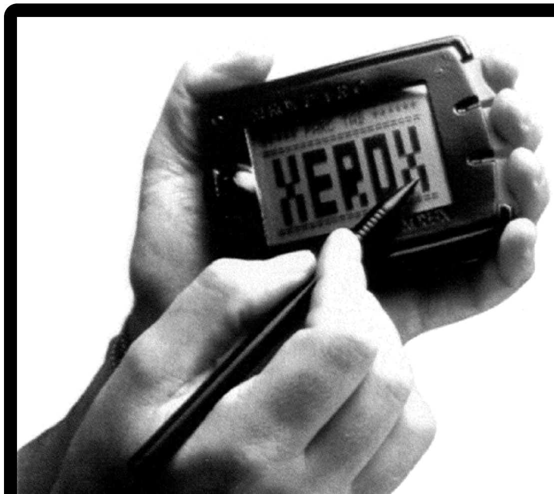

Figure 2.1: ParcTab device (Weiser93)

While the "UbiComp" vision has been influential and prominent for many years, few usable design principles or prototypes have been developed since the original concept was articulated. This work is consistent with the broad vision of Ubiquitous Computing, but is more explicit about the types of hardware and interactions it hopes to support.

Cooperative buildings

More recently, the concept of UbiComp has influenced the field of computer-supported collaborative work (CSCW). The first conference on Cooperative Buildings (Darmstadt, Germany, February, 1998) centered upon integrating technologies with workplace environments. The motivation for this trend was described in the conference proceedings:

> The introduction of information and communication technology has already changed processes and contents of work significantly. However, the design of work environments, especially physical work spaces such as offices and buildings, remained almost unchanged. It is time to reflect these developments in the design of equally dynamic, flexible, and mobile work environments. (Streitz98a)

A representative example of this approach is the "i-LAND" system described in (Streitz98b). This specially constructed room contains a wall-sized, rear-projected display, chairs with tablet computers installed within swiveling armrests, and a bottom-projected tabletop display. Users interact with the devices, and transfer data between them, by using pens and by pointing. The equipment is intended to support group meetings and discussions.

The design of i-LAND, like that of many other computer-augmented environments, implies several common assumptions about computer-augmented interaction:

- that collaboration support systems require specialized, dedicated facilities
- that users in this environment are concerned only with the manipulation of virtual information
- that physical objects are not intrinsic features of collaborative physical environments (no environmental objects appear in the images, aside from a single plant)
- that graphical interaction techniques, with some modifications, are well-suited for multi-user interaction
- that displays must be engaged at distances of 0-10 feet

Collaborative documentation and meeting support may be well served by such facilities, but actions involving physical objects are unknown and irrelevant to them. As a result, it is difficult to imagine how these systems could be employed in environments where physical tasks and objects are central.

2.2 Situated information spaces

The concept of "situated information spaces" proposed by George Fitzmaurice has a somewhat stronger relationship to this work. This work emphasized the importance of physical objects and space in providing a meaningful context for interaction.

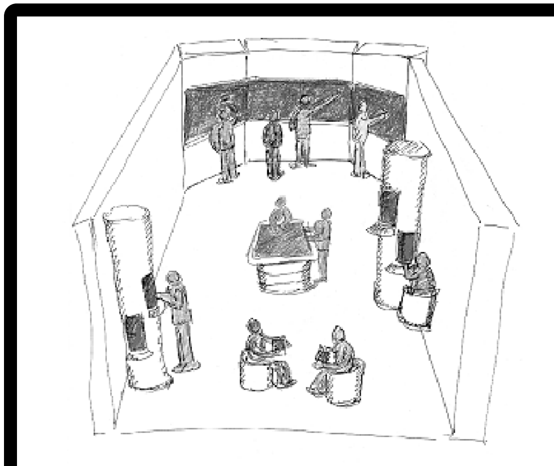

Figure 2.2: "i-LAND" concept sketch (Streitz98b)

Figure 2.3: "i-LAND" "roomware" (Streitz98b)

> Wherever possible, we should look for ways of associating information with physical objects in our environment…. Our goal is to go a step further by grounding and situating the information in a physical context to provide additional understanding of the meaning of the space and to improve user orientation. (Fitzmaurice93)

Situated information spaces were to be enabled by the technology of "spatially aware palmtop computers," handheld displays carried by users as they navigated personal spaces. These devices, typified by the "Chameleon" prototype, were equipped with sensors to detect their positions and orientation. Movement of the device caused the displayed image to scroll over a virtual workspace. The device acted as a "lens" that could reveal a virtual world hidden within the physical.

In addition, proprioceptive displays might be used in conjunction with a distributed system of displays embedded in physical storage structures. Though such a system was never implemented, Fitzmaurice imagined a computer-augmented library with touch-sensitive LCD indicators along the front edges of the shelves: "as we walk through the music section, books on the topic of interest as well as related material will be highlighted by indicator lights…" (Fitzmaurice93)

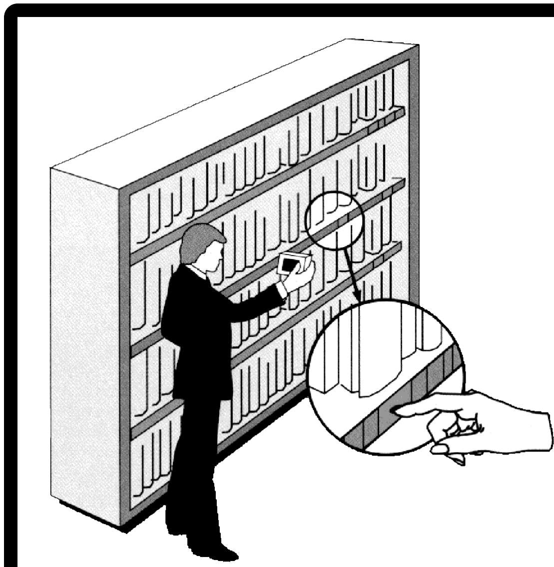

Figure 2.4: Computer-augmented library concept (Fitzmaurice93)

Again, while the broad observations regarding physical context do relate to this work, significant constraints are imposed by the choice of implementation. Use of a small, handheld display requires attention to be focused upon a small screen, and offers no benefit to other users of the same space. Wielding the device continuously also seems incompatible with other concurrent activities. The means to determine the books' positions (for registration with the LCD's) was not addressed, nor were issues of power, networking, data storage, or latency.

2.3 History-enriched digital objects

The concept of automatically embedding representations of history was first proposed in a seminal paper by James Hollan and Will Hill (Hill92). Interestingly, the notion was conceived as a means to facilitate collaboration over digital objects, rather than physical ones:

> The basic idea is to maintain and exploit object-centered interaction histories: Record on computational objects (e.g. documents, menus, spreadsheets, images, email) the events that comprise their use, and then, on future occasions, when the objects are used again, display useful graphical abstractions of the accrued histories as parts of the objects themselves.

Hill and Hollan implemented tools for recording the cumulative time spent editing or reading each line of a document. This data was then represented visually as a simple graphic in the "attribute-mapped scroll bars" alongside the document. This was an elegant solution in that it overlaid the visualization with the area of the screen used for navigation.

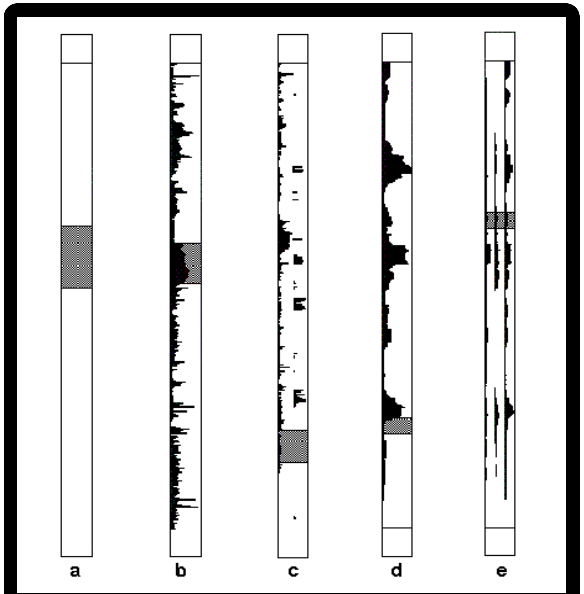

Figure 2.5: Attribute-mapped scroll bars (Hill92)

In addition to describing their implementation, Hill and Hollan pointed out the possibility of inventing new "physics" when generating wear through computation:

> Computation enables the creation of virtual worlds that resemble the real world and allow us to exploit our extensive knowledge of the world in interacting with them… these same techniques also allow us to create virtual worlds that give concrete existence to abstract entities operating according to a physics of our choice. The entities and their physics can be designed to highlight aspects of phenomena not normally available to us but that are important for supporting understanding and task performance.

This point applies equally to "wear" represented in physical environments; use data can indicate information about users, object attributes, modification dates, etc. as well as accumulated use.

2.4 Tangible interfaces


Figure 2.6: MediaBlocks sequencer (Ullmer98)

The work of the Tangible Media Group, founded and led by Professor Hiroshi Ishii, was the immediate context from which this research emerged. "Tangible media," a phrase that signifies both physical embodiment and ease of comprehension, was defined broadly as "the seamless interfacing of physical and digital worlds." (Ishii97) Since its inception, the group has engaged in a series of design projects that have served to demonstrate new concepts, forms, and applications of information technologies. Presented here are some themes that have recurred throughout several different projects.

One theme is the use of physical objects as representations of electronically stored information. Various projects have employed physical objects as "containers, conduits, and controls" for digital information. (Ullmer98) Such systems have the benefits of allowing agile manipulation of data elements, concurrent manipulation by multiple users (Fitzmaurice95), expression of correlations and hierarchies between multiple data elements (Gorbet98), and the learning of motor tasks through object manipulation (Underkoffler98).

Another theme has been the representation of remotely acquired information streams. "inTouch" (Brave97) allows direct and conscious communication between remote parties through the concurrent manipulation of electronically coupled physical objects. Other systems display remote activity or information though a variety of abstract representations: patches of light projected onto walls, rows of spinning pinwheels, ripple patterns on the ceiling of a room. These projects raised many issues about representing remote data displaced from its original context (Wisneski98, Dahley98).

A third area of research has been the degree to which information systems require conscious attention to the display of data. In the ambientRoom and Ambient Fixtures projects (Dahley98), attempts were made to provide "background information," typically live data streams represented through physical phenomena or abstract visualizations. In theory, users could remain unaware of these information sources until sudden changes brought them into the "foreground."

Unfortunately, while a number of novel display devices have been constructed, none were in practice linked to live data streams. Examples of online data sources were suggested—changing stock values, quantity of incoming email messages, activity of a family member in a remote location, etc., but none seemed compelling enough to warrant the effort of implementing a working connection. One problem, it seems, is that most data available online are already well supported by the design and usage context of the personal computer. Conversely, the types of information about which we care most—the activity of people in a different space or time—are simply unavailable online. One goal of my work is to provide information sources to address this problem.

Figure 2.7: inTouch prototype (Brave97)

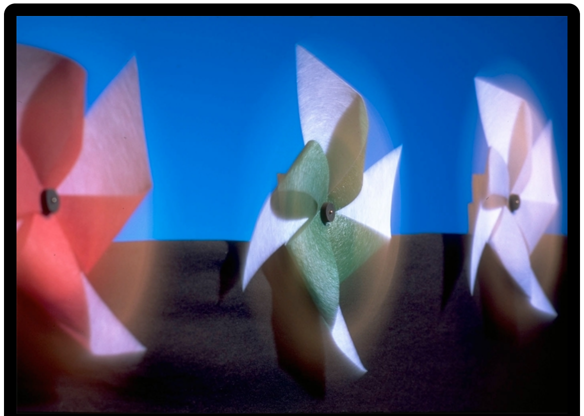

Figure 2.8: "Pinwheels" (Dahley98)

2.5 Artificial intelligence, perception, and emotion

Limitations in the interfaces to computer systems are also being addressed by efforts to grant computers human-like cognitive, perceptual, and emotional abilities. These are mentioned here because many people imagine computers of the future as agents, robots, or anthropomorphic computers that engage users in humanoid dialogue and social behavior. Perhaps the difficulty of achieving these goals should underscore the importance of applying good design practice to existing technologies. In any case, humanoid computers can only benefit from knowledge of human physical actions.

Artificial intelligence

Artificial Intelligence seeks to imbue computers with human-like cognitive abilities through the structured representation of memory and thought processes. While great strides have been made in certain areas—natural language processing, expert systems, and image processing, for example—the general-purpose thinking machine seems at least several decades away, and is perhaps is for our lifetimes a utopian fantasy (Minsky85; Dennett98).

Perceptual user interfaces

The field of Perceptual User Interfaces is a recent outgrowth of computer vision research. It amassed sufficient interest to open its first conference recently (Turk97). "Perceptual" interfaces involve the use of vision, speech recognition, and/or speech synthesis to enable sound- or gesture-based communication with computers. While some systems emphasize passive monitoring of user actions, most are intended as controllers for the familiar GUI workstation.

Affective computing

The focus of Affective Computing (Picard97) is the emotional state of humans and of machines. Various sensing technologies are used to infer information about a user's emotional state. Gesture, posture, speech, and physiological analysis are used to determine the user's level of arousal or confusion when using a computer system. Affective computing has begun to investigate wearable, passive devices for recording emotions as well. For example, the StartleCam project (Healey98) records digital video whenever the user's orientation response exceeds a particular threshold.

3 Design Approach

3.1 Objectives

The approach taken in this work is "research through design." While many useful sensing and display technologies exist, the majority of them require significant effort to adapt to human-computer interaction. Realizing working prototypes that employ these technologies, preferably in forms that can be conveniently reused, is a major goal of this work. A concurrent goal is to propose interaction techniques for this new class of devices. The basic approach has been articulated, but these concepts must be instantiated through specific applications, environments, and types of data. A final goal is to establish conceptual models and maps that help to define the design space for these devices.

3.2 Design Parameters

The possible universe of situated usage history displays is quite large. As an aid to navigation, we propose a set of design parameters that characterize the dimensions by which these devices can differ. This list is intended to motivate the prototypes described in the bulk of this document. After these artifacts have been presented in detail, they will be examined in light of these properties. Of course, the properties listed here are somewhat subjective, and this list is not exhaustive.

Quantity of measurement

The choice of the quantity of measurement to determine "usage" is critical. Generally, the quantities of interest are affected by human activity, but tracking the use of individual objects may be more directly useful than tracking people directly. Thus both human- and object-sensing technologies are of interest. Additionally, measuring devices might count discrete events, such as the number of sheets produced by a printer, or continuous values, such as the overall noise level in a room.

In determining the appropriate quantity of measure, designers should consider

- The utility of the measured data in the task's "action present" (Hill92)
- The cost, complexity, and difficulty of obtaining the measurement
- The extent to which measurement requires changes to existing practice

Physical form

Situated displays may take a variety of physical sizes and forms. Devices can be integrated into large, permanent fixtures like walls or shelves or can be portable devices that are body-worn or carried. Compact, portable usage monitors can also be attached (parasitically) to larger pieces of equipment. Perhaps most interestingly, devices be can small individually, but can act as a large display collectively.

The technology underlying usage history displays need not be complex. For example, imagine a "book monitor" the size of a pack of gum, containing only a battery, capacitor, mercury switch, and a blinking LED. This device could attach to the spine of a book, such that its capacitor would charge whenever the book was tilted into a reading position. The LED would then blink for several hours thereafter, until the capacitor was fully discharged. Thus, one could easily identify which books in a library had been read recently.

The type of data connectivity between devices has both technical and conceptual implications. Remote storage of data may result in unpredictable time lags, while batteries may result in power failure. Only if data, display, and power are local to a device can it be treated as a self-contained, "active" object.

The physical appearance of a device should ideally communicate aspects of its functionality: whether it can be attached or detached, the location of its physical controls, in which direction it should be oriented. If a display represents an attribute of a "referent" object, permitting detachment may lead to ambiguity about the correlation of attribute to object.

Interaction techniques

Interaction, the dialogue between user and machine, must also be designed for the context of use. Users may be active controllers of the interface, casual observers, or may completely ignore or avoid the output of the displays.

In general, all situated usage history displays should have some automated display functionality; they should make the cost of seeking information as low as possible. In this sense, situated displays are like signs posted in public spaces. However, it may be advantageous to instrument existing actions in an attempt to provide more specifically contextualized information.

Quantity of measurement
- type of phenomenon measured
- method of counting
- degree of user adaptation
- criticality of information

Physical form
- physical mobility of displays
- physical aggregation with other displays
- integration with other physical structures
- physical location of stored data

Interaction techniques
- degrees of user involvement
- number of modes

Visualization
- degree of abstraction
- temporal representation
- responsiveness/dynamism
- relativity of comparison

Table 3.1: Design parameters for situated usage history displays

For example, imagine an electronic music synthesizer that recorded the frequency of use of each key, and that represented this frequency by making the keys glow with different levels of brightness. One might program the system to respond to the user by detecting the most recently played note, and illuminating only other notes played in conjunction with that note in the past (as accompaniment or as parts of a chord).

Under some circumstances, it may be useful to add features that allow deeper queries (perhaps into the histories or attributes of objects). In supporting these operations, one must weigh the functionality against the added complexity of moded behavior. Multiple interaction modes should be represented explicitly, and should be triggered only by explicit user actions.

Visualization

The choice of a visual representation of history data is as important as the choice of the data itself. Depending on whether absolute or relative usage is more important, one might choose a numerical or graphical display of quantity. Many techniques from the field of scientific data visualization (Card99) can be appropriated for this purpose.

Time may or may not be represented explicitly. Displays can show either a single value that represents accumulated usage over time (like the time-in-use counters used in forklifts and industrial equipment), accumulated use over a given time period, or the level of use over time (like the scrolling humidity recorders in museums).

The use of dynamism in graphics is a delicate issue. Large moving graphics draw attention and trigger orienting responses (Reeves96), and thus should probably be used sparingly—only, perhaps, on a single display at the focus of attention. Small movements on multiple displays may be tolerable, however; small, synchronized lateral movements may be perceptually similar to the motion of trees in the wind. Dynamic graphics can be very effective at indicating the coupling of a user's physical actions to their graphical representation, a circumstance in which demanding attention is appropriate.

A final note is that color appears to be a very effective filtering criteria in attentional studies (Pashler98), and thus may be a useful technique for visualization of a phenomenon across multiple displays.

4 Gestural Interaction Devices

FishFace and ShakePad

4.1 Introduction

The first interaction devices developed in this research were portable, self-contained, motion sensing devices. "FishFace" detects hand motion through an electric field sensing device called the Lazy Fish; "ShakePad" senses its own tilting and shaking through an accelerometer. Both use an array of color LED's as an output, and similar graphics were implemented on both devices.

While both record "history" in a technical sense (i.e., they save sampled data points over time), they do so on the order of seconds rather than days or years. Rather, the property that makes them interesting for lightweight interaction is the ease by which they can be located, physically acquired, and controlled. Their design is also interesting in that they can be completely self-contained, including their power supplies; thus they both demonstrate technologies that could easily be integrated into the design of other objects.

4.2 FishFace

The FishFace project was enabled by the availability of a particular sensing technology. Joshua Smith, of the Physics and Media Group, had been developing electric field sensing devices for several years. FishFace was the product of his collaboration with the author during the November 1997 course on Tangible Interfaces. In addition to being familiar to Josh, electric field sensing seemed well suited to interfaces involving handheld physical objects.

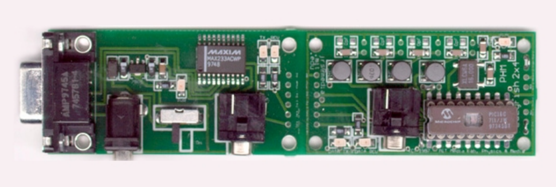

Figure 4.1: Lazy Fish circuit board

Electric Field Sensing

Electric field sensing detects the motion of objects by measuring the changes in electric fields between multiple electrodes in the proximity of the moving object. The oscillating fields are generated by "transmitter" electrodes and are detected by "receiver" electrodes.

These electrodes can be constructed in various sizes and configurations to detect different scales and types of motion. Unlike joysticks, magnetic position tracking devices, and sensing gloves, electric field sensors can detect hand or finger motion without requiring direct contact with the user's skin. This can potentially allow freer motion, greater comfort, improved durability, and lower cost. This freedom from equipment and from wear could be particularly advantageous, for example, in public information kiosks.

Prior to the FishFace project, several people at the Media Lab had employed Josh's electric field sensing modules. In general, the modules were used as non-contact pointing devices for computer displays; users could scroll screens or vary parameters by waving their hands in front of the display. All of these systems were relatively large and stationary. However, Josh had almost completed the "Lazy Fish" module, a self-contained electric field sensing circuit with updated hardware and software. The new module, about 2 x 1 x ½ inches, was much more compact than an earlier version. Josh's goal was to implant the sensors into handheld objects.

FishFace Concept

The field sensing technology was a good match for a different technology that had intrigued the author: the LED (light-emitting diode) matrix. This component is a square plastic tile about 2½ inches square. 64 LED's are embedded in the tile as a grid in 8 rows of 8. Though several students around the Lab had these parts, none seemed to be using them. While their resolution was coarse by computer display standards, they seemed appropriate for a small, simple device driven by a microprocessor. (Ultimately, these devices played a major role in all future projects.)

Several factors made this concept appealing. One was the compactness of the technology. As detection range was approximately equivalent to the electrode dimensions, a device the size of a deck of cards could detect finger motion within a 2" control volume. Thus a user could hold the device in the non-dominant hand while gesturing with the other. Another interesting aspect was the ability to detect motion *prior* to contact—a device might "awaken" as a hand approached it, reacting instantly upon contact. The concept of a single compact device also suggested that unexpected interactions might arise from a multiplicity of such devices.

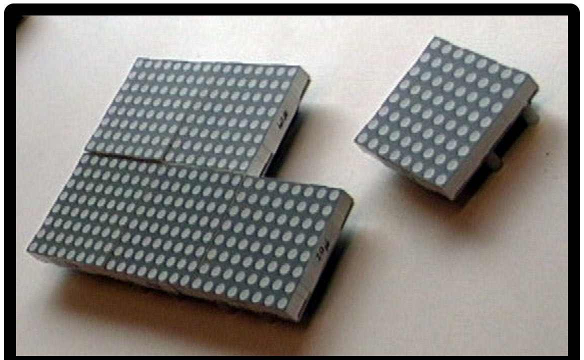

Figure 4.2: LED matrix tiles

The combination of the LED matrix with the “Fish” sensor was sufficiently specific to allow consideration of possible uses and interaction techniques. However, there were several open design questions:

- Would the device be self-contained? Would the LED’s simply provide visual feedback about the detected signal, or would the entire unit act as a controller and display for an external system?

- What types of graphics could be shown on the display? Could treating the display as a window onto a larger, “virtual” space mitigate the resolution limitations? In such a case, the user might move a finger relative to the display to scroll it vertically or horizontally.

- Could the workspace be divided into virtual, invisible “cells,” triggered only when a finger moved into that region? For example, the square could be divided into a 3x3 grid of squares, as are the digits 1 to 9 on a telephone keypad; the entire display could display the numeral corresponding to the selected region.

- Could a pair of these devices be used as a communications medium, like a pair of walkie-talkies? If so, how would local data be mapped to the remote device? Would local data be represented visually?

- What degree of abstraction was appropriate for the graphics on the display? The display was barely capable of representing a single alphabetic character legibly; was there any sense in trying to scroll text messages? Could a simple graphic, e.g. an electronic schematic symbol, be depicted realistically?

As we were able to implement only a few of the many possible configurations, many of these questions remain open.

Original Prototype

The prototype was constructed from two modules: the Lazy Fish sensing board and a new LED matrix control board. Each board had a PIC microprocessor; the two were interconnected by a pair of serial cables. As this was the author’s first exposure to programming the PIC microcontroller, much of the development effort was expended on simply achieving bi-directional communication between the two circuit boards.

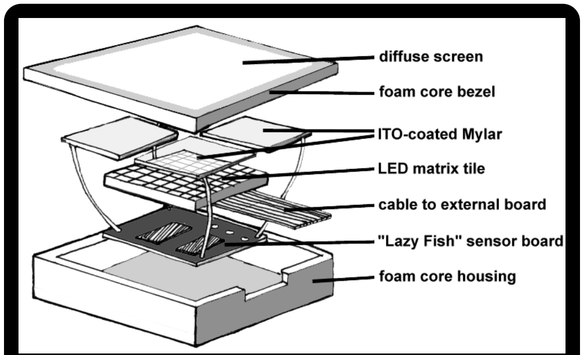

Figure 4.3: Original FishFace assembly diagram

The electrodes were cut from a sheet of Mylar coated with indium tin oxide (ITO), a transparent, conductive material. The transparency allowed the electrodes to be mounted directly atop the LED matrix without obscuring the display. In later versions, the transparent film was replaced by L-shaped cutouts of copper tape mounted around the corners of the display.

Housings were built from a variety of materials—foam core, acrylic, and fused deposition polymer— but were similar in their plain, white appearance. The housings ranged from about three to four inches square in size.

Interaction modes

Several modes of interaction were implemented in the microprocessor code for FishFace. The code for each interaction mode was quite straightforward and differed only subtly from the others. However, the experience of manipulating each one was strikingly different.

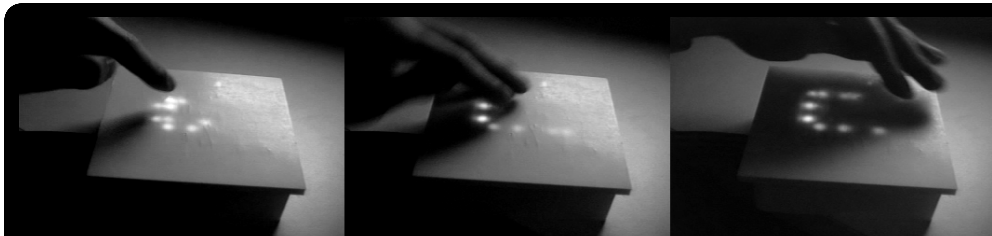

Figure 4.4: Original FishFace displaying "Snake" patterns

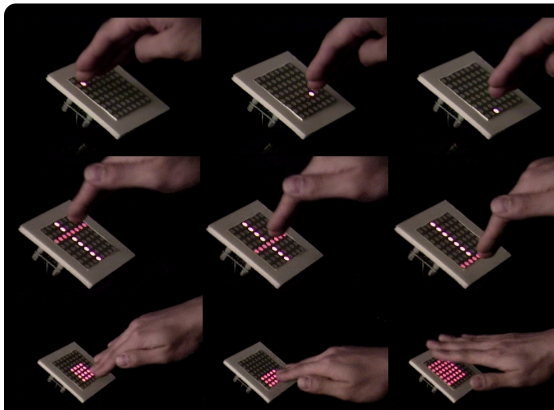

Figure 4.5: FishFace interaction modes "Dot," "Cross," and "Zoom"

In the simplest mode, the sensor's x-y position data was mapped to the position of a single pixel illuminated on the display. The illusion was that the user's fingers were projecting a glowing dot onto the surface of the display, or that an invisible extension of the user's fingers was made visible by the surface. In a slightly different mode, a snakelike chain of pixels follows the user's finger, elongating when moved quickly and collapsing when slowed.

Finally, in a third mode, finger proximity is detected as well as x-y position; a moving square of light appears around the user's finger, expanding as the finger approaches the surface. In this mode, the finger seems to dip into a virtual liquid, the cross-sectional area of the contact plane expanding with increased depth.

Wireless Communications

For an in-class demonstration of the FishFace system, two FishFace devices were linked by a pair of short-range radio-frequency transceivers. The system was implemented hurriedly and with a number of flaws, but for a short time a bi-directional link was operational. The originally self-contained FishFace devices were thus converted to synchronous, portable communication devices, like a pair of walkie-talkies.

To explore this further, one would have to select a model of the mapping of information between the two devices. In a simple case, the paths traced by each user's gestures might be displayed simultaneously on both displays. Another straightforward approach would be to display each user's gesture as a separate color; this would require a bicolor display. Or, rather than displaying each user's motions directly, each user's current position could be used as an endpoint of a line interconnecting the two; thus the line's length would indicate the difference between the user's positions. The inTouch system is a model of another approach; the two displays could indicate a common graphic computed by interpolating between each user's position values.

Performance

The original FishFace display had many technical limitations. The display was monochromatic (red), and thus information could not be communicated through color. Because the display was driven directly from the microprocessor, the available current was barely sufficient to illuminate the display. Room lights had to be dimmed for the display to be easily visible.

Also, reactions to the diffuse overlay were mixed. Some liked the organic appearance of the glowing lines; others found it unnecessarily blurry. The screen was successful in disguising the coarse resolution of the display, however; many assumed that the underlying resolution was much higher than 8x8.

4.3 ShakePad

Several months after the development of FishFace, a similar system called ShakePad was constructed. Like FishFace, ShakePad was a handheld, white, square object with no visible controls. Instead of detecting hand motion through electric field sensing, ShakePad detects its own motion through an accelerometer contained within its enclosure. Tilting or shaking the device triggers changes in the patterns displayed.

Acceleration Sensing

The accelerometer, manufactured by Analog Devices, Inc., is a convenient and versatile sensor. An electronic package about 10mm square, it can be easily integrated into the design of a printed circuit board. Only a few external components are necessary for its operation, and it consumes fairly little power. The two axes of acceleration measurement are in the plane of the device.

Acceleration measurement techniques respond to two different types of motion. Slow, tilting motions are sensed through measurement of the direction of the gravitational acceleration vector *g*. This has a fixed magnitude (9.81 m/s^2) and is directed towards the ground. Thus, if the device is stationary, components of this vector present in the x- or y- axes can be used to compute the angle of inclination. Alternately, if the device is held horizontal but is shaken laterally, acceleration will be detected in the x- and y- axes. As a result, either tilting or shaking the device can produce identical acceleration values.

Graphics

The graphics implemented for ShakePad were based upon those used in FishFace. One difference was the use of bicolor LED matrices; now the colors green and orange were available in addition to red. The use of external LED driver components also made the display much brighter than that of FishFace (which was driven directly from the PIC processor).

In designing graphics for such a system, some filtering must be performed in software to permit accurate control of the dynamic graphics. If shaking is expected to be the dominant mode of device operation, a "high-pass" filter can be used to reduce the effects of gravitational acceleration. Alternately, if accurate tilt measurement is paramount, a "low-pass" filter can be used to attenuate the high-frequency shaking signals. In practice, this technique made the graphics slower to respond, as several values were recorded in order to average the results.

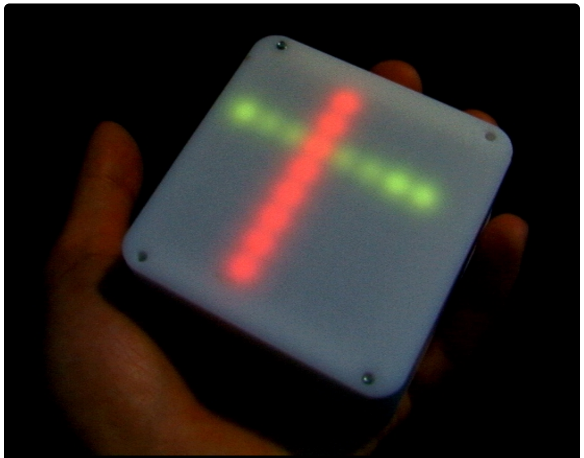

Figure 4.7: ShakePad in Cross mode

The “snake” mode, which featured a collapsing trail of pixels, was implemented for ShakePad. The “head” of the snake, i.e. the most recently recorded pixel, was red, while the snake’s “body” was green. Minimal filtering was performed, such that the snake was very responsive to both tilting and shaking. The inherent noise of the accelerometer was not damped, so the snake always appeared to jump between nearby pixels. The dynamism and color of this interaction made it very satisfying; the snake indeed seemed like an excited creature trapped within the white box.

For the “cross” mode, data was averaged over several readings to reduce the jittering of the previous mode. The display showed a horizontal red line and a vertical green

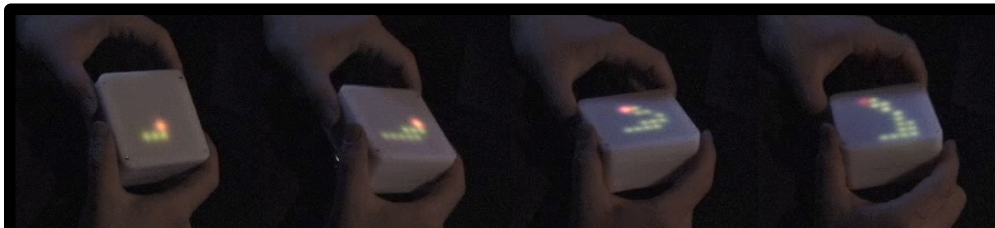

Figure 4.6: ShakePad in Snake mode

line; the position of each depended on the angle of tilt of the display.

4.4 Discussion

LED displays

LED matrices have far lower resolution than common computer displays like liquid-crystal displays (LCD’s), and cathode-ray tubes (CRT’s). Those displays can address over 5000 independent pixels in a square inch, while an LED matrix can manage only 16! Although RGB models are now becoming available, the majority of LED’s can display only red, green, or a combination. LCD’s, of course, can show a full spectrum of colors. In addition, LED matrices consume far more power than LCD’s; this is the main reason they are rarely used in battery-operated equipment.

Nevertheless, LED’s have characteristics that enable fundamentally different types of interaction than the ubiquitous computer displays. Unlike LCD’s, which are reflective or transmissive, LED matrices are emissive. Their high contrast ratio permits viewing at a distance, through diffuse screens, and even through translucent objects.

In addition, the viewing angle of LED matrices is essentially a full hemisphere, while LCD's decrease in brightness when viewed off-axis. The low resolution of LED's, while limiting graphical detail, permits a concurrent reduction in the data that must be sent to the display. An array of LED matrices can easily be driven by a small, inexpensive microprocessor, and a "graphics driver" can be written in a matter of hours. Finally, the displays can be tiled to cover a large physical area without a significant increase in cost; LCD displays cannot be scaled in this way.

The combination of low data rates, high brightness, wide viewing angle, and simplicity give LED's a decided advantage under many circumstances. Most significantly, LED matrices seem well suited for the display of simple information at a distance.

Coupling and responsiveness

User of FishFace and ShakePad experienced the sensation of directly manipulating the glowing pixels. One might expect this to be no more interesting than moving a computer mouse while watching the response of the cursor on-screen. Yet this simple visual feedback was striking—why?

One reason seems to be that the visual feedback was highly responsive—the delay between motion and response was approximately 100 milliseconds. Unless hand motion was very fast, this degree of latency was imperceptible. This made the interactions seem like a physical phenomenon, like a shadow beneath a moving object.

When the positions of graphical objects are continuously controlled by physical gestures, the latency in the control loop seems to affect the perceived nature of the interaction dramatically. (Of course, this is why personal computers have dedicated hardware for handling mouse input.) In both ShakePad and FishFace, the collocation of the display with the interaction space reinforced this illusion. Because the output was closely coupled to the physical input, the moving pixel seemed like a direct extension of the user's body.

Social uses of usage history

The FishFace and ShakePad prototypes are computational devices that can be operated as easily as physical objects can be manipulated. They can be engaged with instantly, and their responses are easily seen and comprehended. Of course, neither of these devices performs any of the functions of traditional computing devices. How might people use them? What roles might they serve in people's lives?

The usage frequency of physical objects can have social and psychological implications. Heavy use may make objects seem more desirable, as popularity may suggest quality, or less so, as repeated use may imply degradation or contamination. Thus, heavily used books, video rentals, tools, or pool cues may be sought out, but heavily used public phones or toilet seats might be avoided. Visible indicators of usage frequency might charge common objects with a new dimension of valence or desirability. Today, the frequent use of communications gadgets confers status; cellular phones with prominent usage displays might become the fetish objects of the week. If every toy in the chest contained a ShakePad that displayed its popularity, what six year-old wouldn't try to grab the brightest one first?

Architecture could also be transformed by the integration of devices that make past contact visible. Imagine an office hallway lined with FishFace devices spaced every twenty feet. The "temperature" of each display could increase with the "body heat" of each passerby, creating a subtle indication of the level of activity in each part of the building. Over time, people would begin to react to this new artifact in their environment. Perhaps they would pause before the displays, swaying back and forth to actively increase their brightness. Or perhaps the ubiquity of these indicators would make less-traveled areas seem especially lifeless and desolate.

Invading people's personal spaces with these devices might have even stronger effects. Consider an airport, bus terminal, or subway car with FishFace devices installed in the benches. The patterns of light could take on the contours of people sitting upon them, only to "evaporate" over time after their users had risen. Would people feel uncomfortable about sitting upon the residual "heat" left by strangers, or would they do so willingly? Would they derive a sense of contact or community from sharing the seat with the shadow of another person, or would they avoid the glowing benches entirely?

5 Augmented Container Storage System

TouchCounters Prototype

5.1 Background

TouchCounters is an integrated system of electronic modules, physical storage containers, and shelving surfaces for the support of collaborative physical work.

TouchCounters evolved from the work environment in which FishFace and ShakePad were developed: the shared physical workspace of the Tangible Media Group. This 20x20 foot area is crammed with workbenches, storage boxes, computers, and raw materials. A revolving set of about a dozen people shares the space. Supplies and materials are stored in 80 to 100 identical plastic containers, labeled inconsistently by attributes such as type of contents, user's name, or project.

The pace of this environment is often rushed and chaotic. Under time pressure and without supervision, people often remove boxes from the shelves and replace them in different positions, thus making it difficult for people to locate specific boxes in a hurry. Reordering of supplies is often done on an emergency basis, and without any tracking of overall usage or efficiency.

Figure 5.1: Storage containers in shared workspace of Tangible Media Group

Initial prototype

Portraying usage frequency seemed an appropriately robust solution for the inherent disorder of this environment. The bright displays of FishFace and ShakePad seemed well suited for use as indicators of box usage, status, and categorization. The original concept was simply to fit each box with a contact sensor, thus enabling its popularity to be seen. This led to the concept of a "touch counter," a "hit counter for physical objects. "

Figure 5.2: Attaching electronic label to storage container

In May 1998, a prototype "augmented container" was developed. A permanent magnet was installed in the lid of the box; a magnetic reed switch on the box could then count the number of opening and closing events. The box was powered by an on-board battery and stored its usage records internally. However, this direction was rejected in favor of a system in which the shelving structure became the substrate for interaction with the system. The shelves would be fitted with multiple "docks" that would provide both power and network connections to the boxes. This would eliminate the need for power supplies on the boxes themselves. During the summer of 1998, this system was implemented.

5.2 Implementation

Labels and containers

The TouchCounters modules were designed to contain the LED matrix, about 2.5x2.5 inches. Each was equipped with a magnet sensor, an accelerometers, and a identification chip etched with a unique 48-bit ID. The electrically erasable microprocessor, the PIC16F84, allowed rapid uploading of new code through a programming connector. Infrared transmitter and receiver modules allowed wireless communication through standard TV remote protocols.

Data was exchanged with the server through a series of connectors on the modules, storage containers, and shelves. Conductive, magnetic snaps were used to link the label modules to the containers. These connectors, sold in the garment industry as snaps for leather purses, carried both power and data signals. The labels could be easily attached and detached with one hand, and tactile and audible “clicks” indicated engagement of the magnets. It was imagined that these labels might serve as generic, reusable usage indicators, perhaps connecting to telephones or to furniture, but those uses have not yet been attempted.

The receptacles on the shelves were wired to a common bus connected to a central server. Standard telephone connectors were used to permit easy expansion. As identical snaps were installed on the containers and the shelves, the act of placing a container upon the shelf activated communication with its label.

Web Server

Originally, the web server was a standard desktop computer running Java™ code. Serial I/O classes enabled the server to read data from the machine’s serial port. To allow remote users to dynamically alter the code executed on this machine, Java’s remote method invocation (RMI) routines were employed. All data was exchanged as ASCII text to facilitate debugging.

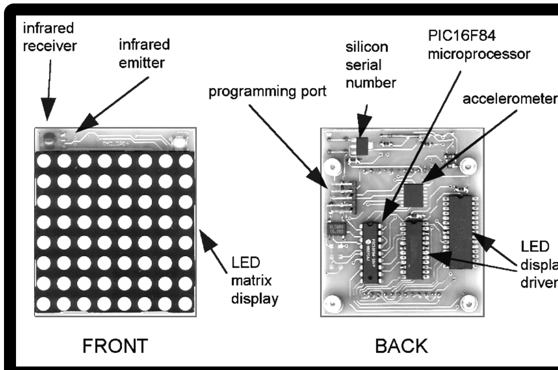

Figure 5.3: Components of label modules


Figure 5.4: Magnetic snap connectors, circled, on containers and shelves

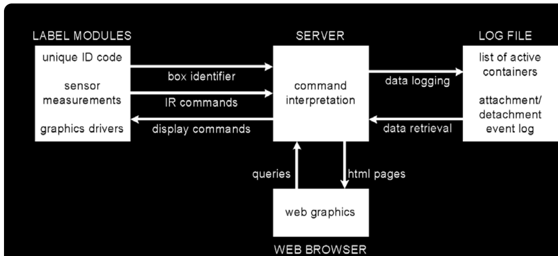

Figure 5.5: Software flow diagram

Later, however, a tiny embedded web server replaced the desktop machine. Its physical compactness suggested the potential for an entirely new type of furniture, designed from the start with integrated sensing and connectivity. This influenced the design of SensePad, discussed in the following section.

In this system, the labels served only as tags and displays, retaining no memory locally; all processing was performed by the server. Descriptions of the containers' contents were manually encoded. Usage correlation was measured both through count information relayed from the labels, and by measuring the time that a box unit was offline and therefore removed from the shelf. Each access event was stored in a continually updated matrix of variables. Likewise, frequency-of-use information was logged in a file available online. A Java applet displayed box status to remote users, and also allowed data to be sent to the labels.

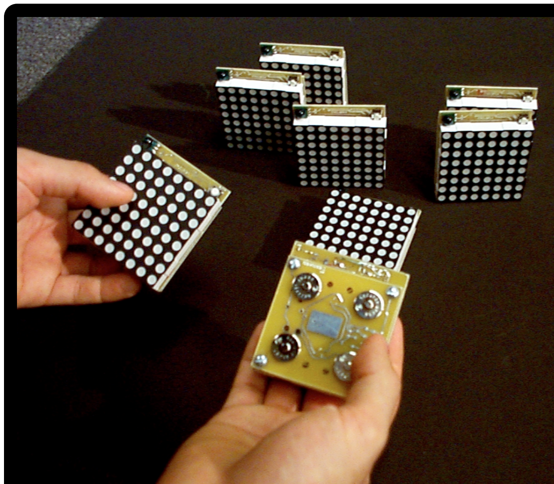

Figure 5.6: TouchCounters label modules

5.3 Interaction techniques

The labels, containers, shelving, and server together supported varying degrees of user involvement in use of the system. These are presented in order of increasing degree of active engagement.

Visualization of usage frequency

The default state of the system is the display of box usage frequency. The label on each box shows a dot pattern that indicates its recent frequency of use; each pixel represents a single use of the container within "short-tem memory," usually a few days. When the entire set of containers is viewed from across a room, the aggregated displays comprise a spatial map of usage frequency.

Figure 5.7: Usage frequency distribution of multiple containers

This provides several types of functionality. As a small fraction of the containers are subject to much heavier use than the average; these "hot spots" can be used as starting points when searching for a commonly used item. In addition, the relative counts can be used to facilitate the manual task of optimizing box placement. For example, the most active ones can be placed at hand- or eye-level. The use indicators also provide an indirect indication of the presence of specific users, as many boxes are associated with individuals. In theory, records of use could be used to prompt replacement of the contents as well.

Visualization of usage correlation

A second mode of operation is the display of usage correlation. Through analysis of the record of box accesses, the system generates a matrix of values that represent the correlation of use between boxes—how many times they have been removed within one minute of each other.

While usage frequency is displayed persistently, this correlation display appears only for the few seconds following the removal of a box from the shelf. (After this time the labels return to the frequency display mode.) To distinguish this mode from the usage frequency display mode, both color and graphical patterns were changed. The normally green graphics changed to red, and the pattern changed from a “solid fill” to a pseudo-random pattern of pixels.

As two or more containers were often used in combination, the correlation display facilitates the rapid location of several related items on the shelf. In this way, the search for a related item can be accelerated by looking at the brightest displays in the area. This data can also be used to improve container placement, as strongly related boxes can be positioned near each other.

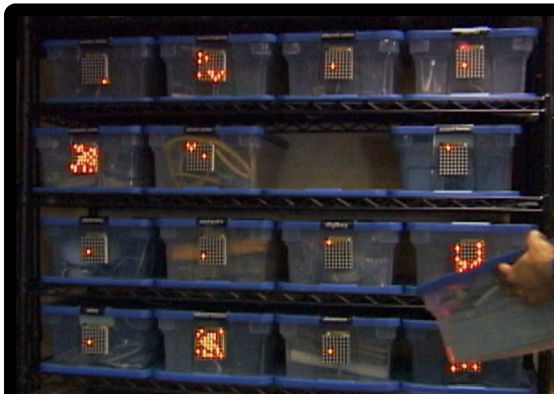

Figure 5.8: Correlation of usage to container just removed

Direct annotation

In addition to displaying automatically recorded data, the system allowed users to explicitly annotate the boxes. By pointing an infrared remote control at the boxes, users could attach symbols or “glyphs” that indicated common associations between several containers. Holding a button on the remote transmitted repeated bursts of data, so multiple boxes could be rapidly labeled by “dragging” the controller across physical space.

The categorization labels were imagined to be arbitrary attributes ascribed by the user, as are the colored “labels” in the Macintosh™ OS. For example, users could indicate associations of boxes with specific users or specific steps in a project. Alternately, a group of users could agree upon certain symbols as indications of the state of completion of various prototypes.

Potentially, users might use other infrared-enabled devices to interact with the labels. Users with personal digital assistants could attach electronic “notes” indicating that an item had been borrowed. Alternatively, users could wear infrared-emitting name tags (Resnick98) to recall personal settings for the system

Figure 5.9: Labeling individual containers with infrared remote control

Online status display

Finally, users could view the state of the entire system remotely by opening a web page with an embedded Java applet. As a demonstration of the bi-directional interaction, users could also click on these web-based images; this triggered changes to the physical displays.

This was a trivial example of how a conventional graphical interface might be used in conjunction with the distributed displays of the TouchCounters system. The GUI can more easily support complex, textual queries, while the physical visualizations would support the frequent and rapid “lightweight” interactions. The GUI’s separation from the physical objects is tolerable or even advantageous in some circumstances; a remote supplier of parts or an inventory manager might access the system to analyze patterns of consumption.

However, the most interesting opportunities seem to arise when fundamentally digital operations can be performed on physical objects. If objects were labeled with metadata describing their contents, users, times of use, and so on, they might be digitally indexed, searched, and located *in situ.* This is a new type of functionality that centralized graphical interfaces simply cannot provide.

Figure 5.10: “Dragging” remote control across containers to label multiple objects

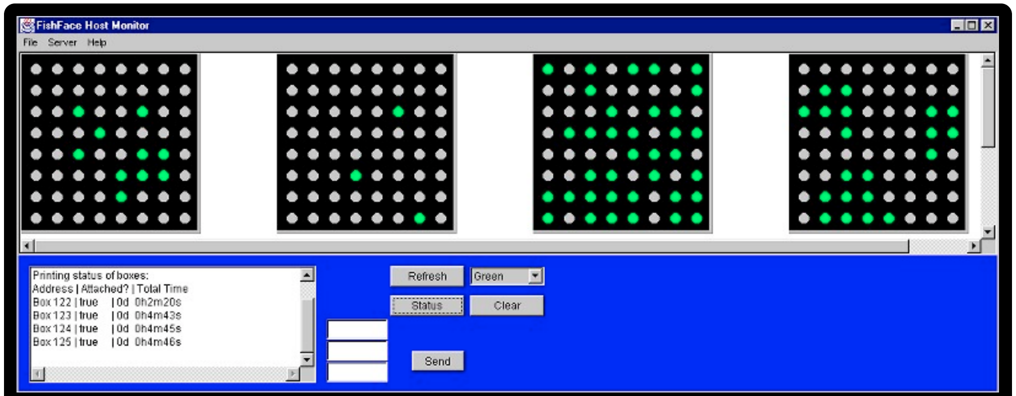

Figure 5.11: Applet indicating status of physical containers on web page

5.4 Discussion

User feedback

Some qualitative feedback about the system was obtained from members of the group. While no substitute for formal testing, this did provide some guidance for design. Sixteen containers were fitted with the labels and networking hardware; the remainder were left unchanged.

Generally, users responded positively to the display graphics, and expressed interest in the complete implementation of the system. In particular, users liked the visual feedback that accompanied labeling with the remote control. Some users asked how messages could be left on the containers themselves; for example, a note that an item had been borrowed from a box. Others asked whether the users of each box could be identified, such that personalized settings or histories could be retrieved.

Interaction at range

The TouchCounters displays were bright enough to be seen from across a room. Three types of viewing could be distinguished, which correlate approximately to three scales of distance between user and display. From 10-20 feet away, the glowing of the displays simply indicated that the system was present and operational. At a range of 5-10 feet, patterns on the multiple displays could be compared. Within three feet of the boxes, the printed container labels could be examined, each pixel on the display could be seen, and the box could be accessed physically.

These levels of scale had implications for the design of the display graphics. The visibility of the displays made it important to prevent them from becoming distracting. In particular, animated graphics were not used as they attracted attention from all points in the room. This made the switch from usage frequency mode to usage correlation mode very noticeable.

Figure 5.12: Viewing TouchCounters at long-, medium- and close-range

wire	solder	scissors	velcro
serial cables	ethernet	serial adapters	foam
fabric	magnets	abrasives	adhesives
paint	markers	tape	monitor cables

Table 5.1: Materials tracked while testing TouchCounters

Different representations of quantity were used in the usage frequency and correlation modes. Both represented only a single value that represented accumulated use, but they suggested different meanings. The "scatter pattern" was intended to appear random, but many interpreted the pattern as a two-dimensional encoding of information. Others complained that the "fill" pattern should fill from the bottom, as a liquid fills a glass. More sophisticated indications of usage history could be represented on the matrix displays, such as a histogram that indicated the temporal distribution of usage events. (See the Network Meter in Appendix A.)

Because all of the individual displays showed the same type of data, they acted in concert as a single, room-filling "meta-display." As persistent fixtures of the environment, the usage indicators might fade into the "background" of one's awareness, noticed only if sudden changes occurred (Ishii97).

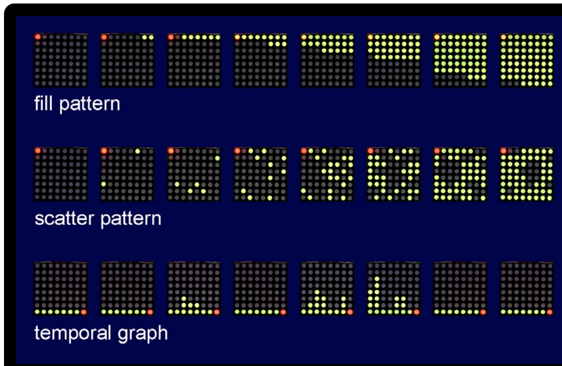

Figure 5.13: Graphical representations of usage

Conceptual models and new technologies

The expectations of any new technology are strongly influenced by prior experience with other technologies. When faced with new developments, analogies to familiar precedents are employed to exploit existing conceptual models. New devices are understood in terms of parallels to older or simpler devices. Thus, designers of radically new interfaces must reflect upon which assumptions engendered by conceptual association must be preserved, and which can be discarded.

Imagine the task of explaining the modern personal computer to a person who missed the last 50 years. One might make reference to a variety of somewhat older technologies:

- television, as a source of dynamic representational images viewed on a screen
- telephones, as a means for synchronous interpersonal communication that transcends distance
- light bulbs, as a consumer of energy and a source of light and heat
- automobiles, as complex machines requiring skill to operate, as means of personal empowerment
- books, as repositories of information
- record players, an automated mechanism for streaming media playback
- human beings, as information input and output systems that make use of language, memory and cognition

To extend the use of TouchCounters to uses beyond storage management, one might attempt to characterize them through analogy to other physical devices. In one sense, TouchCounters can be considered tags or labels, as they are unique identifiers that attach directly to their physical referents (although this operation was not used as part of the interface). They are augmentations to existing objects rather than tools in themselves. However, while a single module indicates the use of an individual object, collectively they portray the activity level of an environment. Also, they are a mechanism for sharing group knowledge in a public space. In these senses, they seem like signs. But unlike traditional signs, TouchCounters show dynamic information. This suggests that that they are indicators, embedded displays of dynamic variables often viewed in parallel.

The point is that when physical devices are augmented by computational mechanisms—memory, communications, procedural behavior—the scope of potentially available uses and functions suddenly explodes. Thus, designs for lightweight interactions must weigh each additional level of functionality against its additional disruption of existing conceptual models. Often the challenge concerns not sophistication but simplicity; the difficulty lies in identifying the few features that can be both useful and easily understood.

	SIMILARITIES	DIFFERENCES
Labels	Describe specific physical objects	Mainly static; edited manually
Signs	Describe physical environments; can be viewed at a distance	Static; not associated with physical objects
Indicators	Show dynamic information, can view several concurrently	Not associated with physical objects, fixed mapping

Table 5.2: Comparison of TouchCounters to other types of physical indicators

6 Object Identification and Tracking Surfaces

SensePad and IDPad Prototypes

6.1 Introduction

While the TouchCounters system was based on an existing system of physical storage, its physical design was a hodge-podge of appendages and modifications to existing equipment. Labels, connectors, and docks clung like barnacles to the shelves and plastic containers, while computers huddled nearby, tethered by many umbilical cables. This was appropriate for the demonstrating the various interaction concepts, but for the deployment of usage history displays on a broader scale.

The SensePad and IDPad projects were attempts to integrate sensing and display technologies into the design of new types of furniture, such that they could be incorporated into various environments with minimal disruption. This goal had several components: (1) to make the sensing technologies more reliable and robust, (2) to reduce the size and the cost of the labels used to identify the objects, and (3) to develop modular physical platforms that could be integrated into tables, desks, shelves, and other furniture upon which objects often reside.

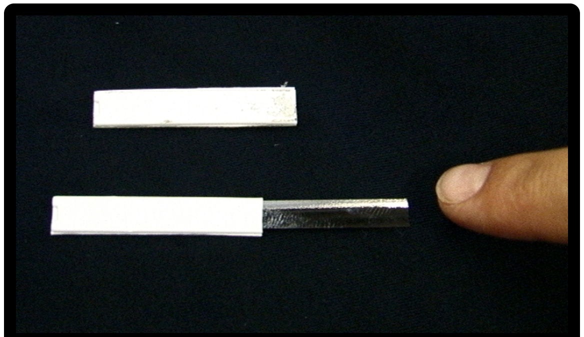

Figure 6.1: Electromagnetic antitheft tag with resonant filament exposed

Tagging Technologies

Broadly defined, a tag is a device designed to facilitate the automated identification of a physical object. The majority of tags are inexpensive "labels" that are affixed to various objects, and which are then identified by "readers." A wide variety of tagging devices exist, but in this section we concern ourselves only with wireless identification tags—those that can be identified without direct contact through interrogation with electromagnetic signals.

Two of the most popular tag technologies are printed optical tags, such as bar codes, and radio-frequency identification (RFID) tags. Optical tags require a direct line of sight between an optical sensor or imaging device and the surface of the tag. While optical tags can be made very inexpensively, they pose difficulties stemming from occlusion, surface contamination, and registration. The most common RFID tags are batteryless electronic chips that contain a simple identification number. These tags can be "read" by exposing them to an oscillating electromagnetic field. While drawing power from this field, these tags then transmit an identification number by modulating the field. While these tags are compact, convenient, and offer both read and write capabilities, their high cost has prevented them from replacing optical tags. (Typically, RFID tags cost from $0.50 to $10.)

Another common type of identification tag is the electromagnetic anti-theft tag, of the type manufactured by Sensormatic and Checkpoint. These tags allow stolen goods to be detected upon their removal from a store's premises. While these tags are very inexpensive ($0.01), they have essentially no memory capacity and thus cannot be used to distinguish amongst multiple items.

"Materials Tags"

Rich Fletcher, a doctoral candidate in the Physics and Media Group, was developing ultra-low-cost wireless tagging technologies. Rather than storing information on semiconductor memories, these tags would store unique patterns in soft magnetic materials; thus they were called "materials tags" (Fletcher96). The spectral response of these tags would be modulated by the magnetic materials. If realized, materials tags could dramatically expand the application of tagging technologies; for example, every item in a supermarket could bear a unique tag. Clearly, a variety of new interface needs would be created.

As a short-term substitute for the eventually forthcoming materials tags, EAS (electronic article surveillance) tags manufactured by Sensormatic could be detected by tuning the device to 58kHz. The frequency of these tags could be increased up to 70kHz by opening their plastic packages, removing the metal strips inside, approximately trimming their lengths with clippers, reinserting them into their plastic shells, and then graphing their frequency response using a PC-based application. A limitation of this approach was that the tags did not exhibit a single resonant frequency but instead "double peaks" a few kHz apart. This limited the number of tags differentiable in the 58-70kHz range to perhaps 20.

Tag Reader

The primary product of Rich's efforts was a swept-frequency tag reader optimized for materials characterization. Through measurement of the frequency and strength of the reflected signal, the reader could determine the tags' proximity to the antenna, as well as their identities. Unlike commercial tag readers, this device could be tuned to detect tags at frequencies from 0 to 300 kHz. In addition it was capable of reading

Figure 6.2: Frequency response of several LC tags; distinct "peaks" are visible

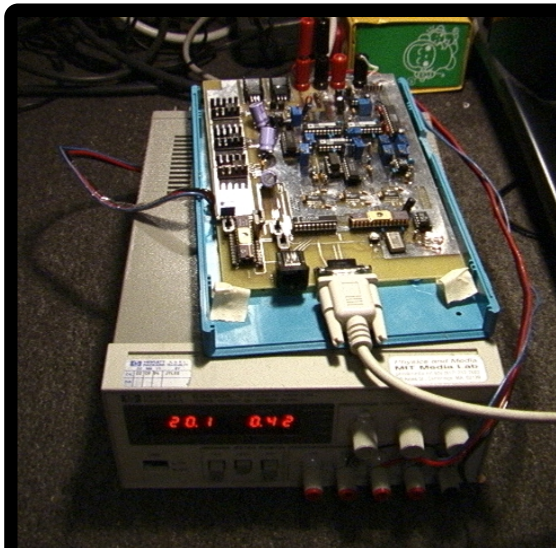

Figure 6.3: Swept-frequency tag reader atop power supply

both magnitude and phase of the detected signals, a requirement for reading materials tags. The tag reader's physical form was a 6x8" circuit board in a blue plastic enclosure. It was usually connected to a 20" diameter loop of coiled wire in a plane perpendicular to the tags. The same coil would be used to alternately "excite" any tags in its proximity and to detect the oscillations of any tags present.

6.2 SensePad

SensePad was intended to be a device that displayed graphical patterns directly beneath objects placed onto its surface. The system was designed to both recognize and track the positions of objects bearing wireless identification tags. Unfortunately, despite a prolonged development period of over 18 months, its intended features were never completely realized. The IDPad project, which arose somewhat serendipitously, later achieved some of the goals of SensePad.

SensePad Concept

From the spring of 1998, the author collaborated with Rich Fletcher to develop a visual interface to Rich's tag reader. The interface was envisioned as an "interactive surface" that would react to small, tagged objects placed upon it. After some discussion, it was agreed that the device would be an oblong platform with LED's upon its top face; it would reside on a tabletop and would exchange information with a desktop computer. The graphical display would indicate (1) the position of each object along the length of the strip, (2) the proximity of each object to the surface of the strip, and (3) a symbol that distinguished each object from the others.

A real-world application had emerged in parallel with the development of the SensePad concept. "Medication non-compliance," estimated to cost $1 billion per year, has been addressed by a wide variety of technologies—electronic bottle caps, voicemail messages, pill box alarms, pill removal sensors, etc. (Cramer91) These systems provided either the function of reminding patients to take their medication, or a means to record whether medication had actually been taken. Becton-Dickinson, a Media Laboratory sponsor, expressed interest in a device that could record the usage of medications in the home.

As a visual interface to information about tagged objects, an interactive surface like SensePad seemed well suited to serving this function. As a reminding device, SensePad could illuminate icons beneath medicine bottles upon its surface to indicate

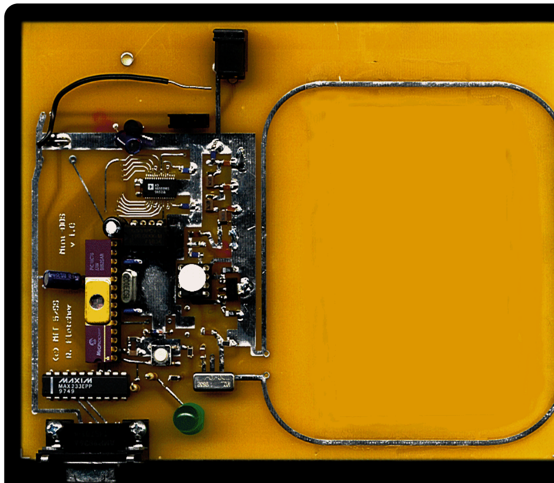

Figure 6.4: "Mini" tag reader, developed later; antenna coil requires only one loop

which ones were due for consumption. As a recorder, SensePad could detect when each bottle was removed (and its contents, presumably, consumed). Furthermore, SensePad could be sized to fit into a user's home medicine cabinet.

Additional functionality specific to medication monitoring could be added as well. Information about a patient's medication regimen could be downloaded from a remote database or perhaps stored in the medicine bottles themselves. Potentially, the system could also indicate warnings if a contra-indicated medications were taken together.

First Prototype

The development of SensePad was a protracted and somewhat inconclusive effort. This process is documented here for readers interested in the underlying technology; others can skip to the next section.

In keeping with the above concept, the initial prototype was a 2.5"x12" pad with a monochrome red 8x32 pixel display (comprised of 4 LED matrices). As with FishFace, the graphics would be coarse but bright and responsive. The detection of tag position would be achieved by measuring signal values at each of 4 coils along the length of the strip.

The device would connect to Rich's tag reader through an external cable. Because Rich's device was designed to use only a single RF coil, a switching circuit on the new device would be necessary to allow communication with multiple coils. In this prototype, switching between coils would be performed by a series of electromechanical relays controlled by a microprocessor on the device. In addition, a separate microprocessor (the Microchip PIC16C73A used in FishFace) would be used to generate graphics on each of the LED matrices. An additional processor of the same type would be configured as a "master" that would control the four "slaves." Finally, a small beeper would provide auditory feedback.

The circuit board for the first prototype was designed in mid-April 1998; components were installed and test software written over the next few weeks. Unfortunately, a number of problems made the board virtually unusable. To illustrate the technical difficulties that characterize this type of work, these will be described in some detail.

Figure 6.5: Medicine bottles with electromagnetic tags attached

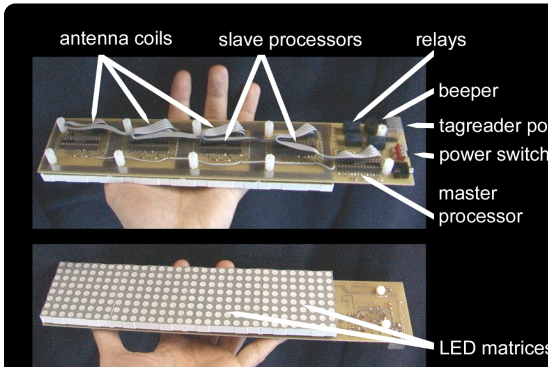

Figure 6.6: Components of initial SensePad prototype

The first problem was that the tag reader could not successfully interrogate the tags. The coils, both smaller and with fewer turns than the original coil, could not excite tags placed near their edges. Thus, unless tags were placed in the center of each square, they could not be detected at all. In addition, the proximity of the analog antenna traces to the digital data and graphics traces caused noise on the tag signals.

Electromagnetic interference was another problem. Rapid switching of the relays caused inductive interference that frequently reset the microprocessors. No mechanism was available for resetting an individual processor; the entire board had to be shut down if the code on one failed. The switching also created a loud, unsettling clicking noise.

Another set of problems concerned the multi-processor architecture. Revising the graphics required the tedious removal, erasure, programming, and replacement of five microprocessors. As the slaves handled communications in software (rather than by hardware UART), command reception could not occur in parallel with the display of graphics; this made the display flicker visibly.

Difficulties also plagued the exchange of data between the master, the slaves, and the processors on the external tag reader board. These commands were series of ASCII characters sent via the RS-232 serial protocol. As the tag reader's software had been successfully written and debugged earlier, it was taken to be immutable. However, that software was optimized for materials characterization, rather than instantaneous response. Several command variables had to be sent to the reader between coil-switching operations; this slowed its operation.

Other problems included the brightness of the display, as dim as that of the original FishFace, and various difficulties in assembly of the board. Clearly, a major redesign was in order.

Second Prototype

A number of improvements were made for the next version. The first addressed the problem of read range. Fortunately, the tag reader had separate terminal connections for coil excitation and detection—these functions could be performed by separate coils if they were close to each other. On the revised board, a single, large transmit coil was created that would surround the four original coils. Its field would thus ex-

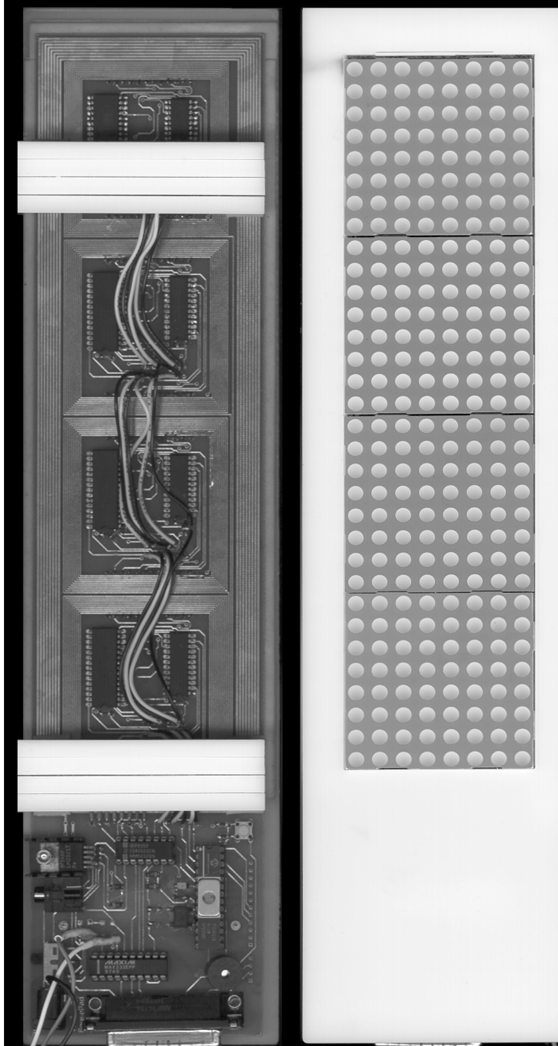

Figure 6.7 Second SensePad prototype, top and bottom surfaces

cite all tags on the SensePad surface, while the original coils could then be used for detection.

Separation of the transmit and receive coils offered another advantage. While current through the transmit coil had to be high, the current induced in the receive coils was low. As the transmit signals no longer were switched, the relays could be replaced by an analog multiplexer chip with a relatively high impedance. This device would switch the received signals electronically, thus eliminating the electrical and acoustic noise caused by the relay inductors.

To reduce the complexity, cost, and instability of the five-processor configuration, each slave was replaced by a pair of LED driver chips. These devices could be chained sequentially, thus allowing a single processor to control the entire array. The drivers were capable of sourcing a high, constant level of current; this would allow the displays to be uniformly bright. They also supported more output lines than did the PIC microprocessors. This allowed 24-pin, bicolor LED matrices to replace the original monochrome displays. Each pixel could appear green, red, or a yellow-orange combination of the two.

In addition, the “brains” of the device were now split across two circuit boards. Communication with the tag reader board would be handled by a microprocessor mounted directly on that board; the chip on SensePad itself would handle only graphics and coil switching. This change required manufacturing a new tag reader board as well as a new SensePad. The new circuit board designs were dispatched for manufacture in mid-July, but returned with some errors; a corrected version was first assembled in August.

The device could now measure the response of electromagnetic tags placed upon each of the four coils. The reflected signal was strongest when tags were placed at the center of each coil, and diminished as they moved towards the edges. Thus, the position of a tag could be approximately computed by weighting the signal strength of a tag at each coil by a number representing that coil’s position. The proximity of the tag could be computed by simply summing the strength values at each coil. This technique had the limitation that position could not be accurately determined if the tag was placed at the extreme ends of the pad, past the centers of the first and last coils. This meant that only three-fourths of the display area could not be used for sensing. Nevertheless, the simple weighting scheme allowed reasonably accurate position tracking.

At last, the board's sensing hardware was sufficiently operational to allow experimentation with the graphical output. The control processor could now compute the identity, position, and proximity of each tag. For demonstration purposes, it was decided that a graphical patterns for to three to five tagged objects would be sufficient to indicate the device's functionality. It seemed natural to employ a graphical symbol or "sprite" that would follow the position of the tagged object, as a shadow follows a physical object. The sprite for each tag could then have a different shape. To indicate proximity, the shape would "grow" as the tag approached the SensePad surface.

The shapes had to be visually distinct, and had to be larger than the footprint of a medicine bottle but no larger than the width of the display. Since orientation information was not measured, only radially symmetric patterns would be used. (This also allowed a reduction in memory space consumed by each graphic.) A rounded, "disc" shape, a square, and an expanding spiral shape were the first shapes chosen. For each type, eight variations were generated in sizes ranging from 4x4 to 8x8 pixels. In summary, the following mappings were used to indicate the object's parameters:

- object's identity: shape and color of symbol
- object's location: position of symbol
- object's proximity: size of symbol

At the inception of this project, the graphics were imagined as shadows that would closely follow the motion of the tagged objects. Of course, such an effect would depend upon the rapid tracking of object motion. Unfortunately, as the tag reader took 2-3 seconds to measure the position of each tag; the graphics lagged the objects' motion very noticeably. For demonstration purposes, it was necessary to move the objects very slowly to create the illusion of coupling. It was clear that improving the speed of the tag reader was essential.

This system was demonstrated to various sponsors at the Things That Think Consortium meeting in September 1998. During the following week, a meeting was held with several executives from Becton-Dickinson Corporation. They felt that the device was too small to accommodate the medication collections of some patients, and suggested that, to be effective as a medication monitor, the device be redesigned to hold at least 20 bottles simultaneously.

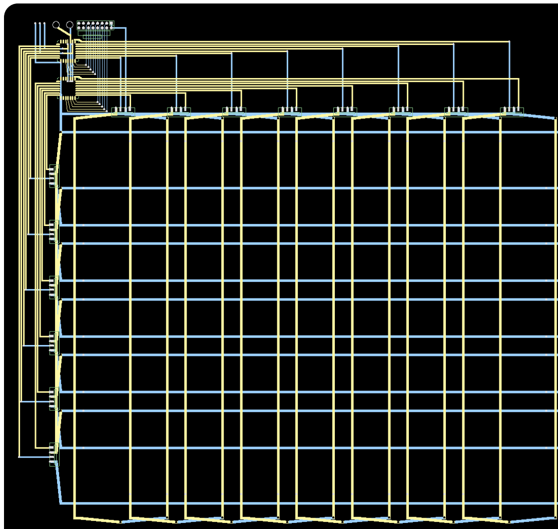

Figure 6.8: Antenna coil for third prototype

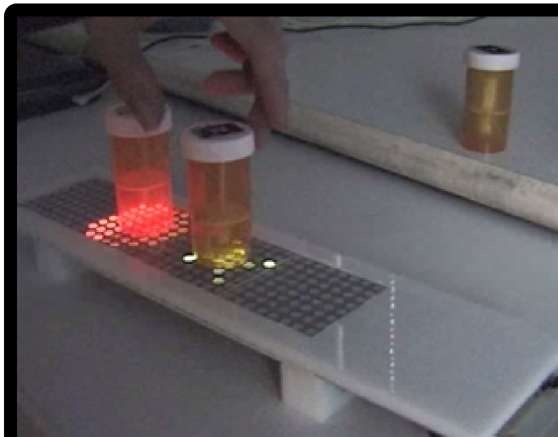

Figure 6.9: Second SensePad prototype tracking two medicine bottles

Third Prototype

Many changes were made in the next version. First, the display surface was expanded from four to twelve tiles, arranged in a 3x4 pattern. This size could accommodate 20 medicine bottles comfortably. The display area was now approximately equal to the printable area of a 8½x11" sheet of paper; future applications were envisioned involving printed templates backlit by the display. Because many more pixels would have to be updated per second, a faster processor was needed to drive the display. The Hitachi SH-2, a 32-bit processor, was chosen for its high speed, low cost, and convenient development tools.

Meanwhile, in parallel with the development of this prototype, Rich had constructed a smaller, faster, and less expensive tag reader. This so-called "mini" tag reader was a much simpler, faster, and lower-power circuit for use with high-frequency "LC" tags in the 5-15 MHz frequency range. Because of the higher frequency, this reader required only a single loop of wire for excitation and detection; this allowed the entire device, including antenna, to be constructed on a single 4x6" PCB.

The simple antenna format enabled the layout of a two-dimensional tracking array. A new PCB was designed with 6 horizontal and 8 vertical loops. By testing each of these overlapping coils in sequence, the 2D position of a tagged object could be determined. Two multiplexers were used to switch the high and low sides of the signal.

This tracking array was designed to fit on top of, rather than beneath, the display surface. Although it could not transparent, it was designed such that a large rectangular window could be cut from the center of the PCB to allow the display to show through. The missing traces would be replaced by thin wires stretched taut across the window in both horizontal and vertical orientations. The display would then be covered with a diffuse screen to conceal the wires beneath.

The LC tags, which consist of an inductor and a capacitor in series, could be designed in various form factors to specific frequency values. The inductor for this type of tag could be made from a spiral trace on a printed circuit board or from wire finely wound into a loop, while the capacitor was a standard electronic component. The LC tags had sharper and more stable frequency "peaks" than the electromagnetic tags. This meant that more unique tags could occupy a given frequency range.

Figure 6.10: "LC" tag made from coiled wire and capacitor

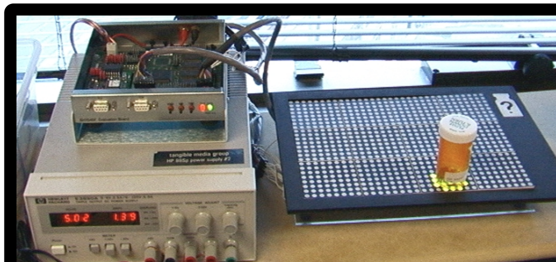

Figure 6.11: Expanded display for third prototype; power supply and control board also visible

The new display circuit was bright and flicker-free. The SH-2 also provided a much more professional development environment. The sensing circuit, however, had some problems. The multiplexers' ON-state resistance seemed to interfere with the reading of the signal. In addition, the mini tag reader lacked sufficient transmit power to drive these coils; it would have to be redesigned with a high-power output stage.

Conclusion

While these problems were not intractable, due to timing complications, they were never completely resolved. The author hopes that another party may be able to complete the design in the near future. Nevertheless, a functional wireless tag tracking system was eventually constructed—the "IDPad" system described in the following section.

6.3 IDPad

The IDPad is an alternative tag-tracking interface device that was originally built as an test platform for a revised version of TouchCounters. After its construction, however, it seemed useful as a platform for direct manipulation of objects rather than usage history tracking. Nevertheless, the technologies and design issues are discussed for their relevance to the broader work.

IDPad Concept

IDPad was an outgrowth of the TouchCounters project, rather than of SensePad. The design of the TouchCounters system had required different components on the label modules, containers, and docks. Many of the connection points were unreliable, which made continued operation of the system very difficult. A second-generation system was conceived in which the displays and sensors would be embedded in the shelves, while the boxes would simply bear RFID tags. No electrical connections would be engaged or disengaged during operation, thus improving reliability.

Each shelf in the storage rack would be augmented with the following components:

- A commercial RFID tag reader board
- Four large antenna coils mounted along each shelf's surface
- An RFID reader controller board to link the tag reader with each of the four antennae
- Four 4x8 LED matrix displays mounted along the front edges of the shelf
- A driver board for controlling the LED matrices

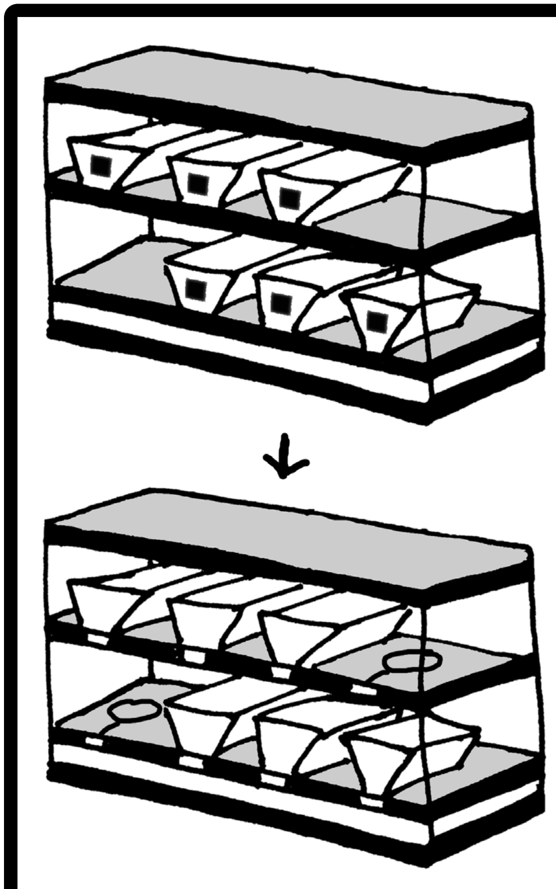

Figure 6.12: Concept for revised TouchCounters system: displays integrated into shelving

As soon as a box was placed upon the shelf, the RFID controller would identify the tag attached to its underside, and would then send this information to a central server over a serial bus. The server would then generate a graphic pertaining to that particular object and send it to the display controller on that shelf. The display beneath the newly inserted box would then illuminate with the graphic. These operations were intended to occur "instantly" (within 200-300 milliseconds).

RFID Technology

While the LC and electromagnetic tags used in SensePad prototypes are simple, single-frequency resonators, RFID tags are complete semiconductor devices with onboard memory and logic. RFID tags draw power from radio waves oscillating at a given frequency. Communication between tag reader and tag is based on modulation of this field; any arbitrary sequence of bits can be sent in this fashion. Thus, RFID tags can have far more sophisticated functionality than the simple antenna-like materials tags. Far more objects can be uniquely identified, as they have rewriteable memory of 256 bits or more.

RFID is an inherently digital, while resonant tag sensing is analog. The signals of resonant tags fall off slowly as they move farther from a sensing coil, while RFID signals simply stop at a certain threshold. An important implication is that the position of RFID tags cannot be tracked continuously, but only at the discrete locations of multiple coils. This constraint was fully acceptable for the TouchCounters application, as placing the four containers at discrete positions was not inconvenient.

Figure 6.13: Card-, coin-, wedge-, and pill-sized RFID tags

Implementation

Around March 1999, the RFID reader was obtained, and the RFID controller and display boards were designed. In order to test these components prior to their installation, an enclosure was fabricated using a laser cutter. This enclosure was a long, narrow pad sized to contain the four antenna coils and four displays.

The author was determined to avoid the technical failures of the SensePad project, especially those involving tag reading technology. Thus the tag and reader devices were chosen from proven, commercially available components.

Tag reader. Despite the relative ubiquity of RFID technology, it was difficult to locate a tag reader vendor that sold unit quantities at reasonable prices. Eventually, a small design company was located in the UK; it had developed a PIC-based tag reader that could read and write Philips tags in the 125kHz range. The device was conveniently postage-stamp sized, and could be controlled by TTL-level RS-232 signals. The company, IB Technology, provided invaluable technical support over the telephone.

Tags. Tags were also surprisingly difficult to obtain. Strangely, it was impossible to obtain tags from Philips, the manufacturer. Dozens of calls, faxes, and emails eventually unearthed a few vendors with short lead times and small quantities. Several tags were obtained in several form factors. One was an opaque, white, and the size of a credit card; another was thin, flat, and nickel-sized; yet another was a tiny, pill-shaped glass capsule intended for animal injection. The type primarily used, however, was a black plastic wedge 12x6x3mm in size.

A significant limitation of the tag protocol was the inability to communicate with multiple tags concurrently—a feature known as "anti-collision technology." If any two tags were within range of a single coil, all reading or writing was disabled. Thus, the four-coil surface could identify at most one tag per coil.

Antenna coils. After brief experiments with hand-wound coils of enameled wire, the antenna coils were designed on a PCB layout program. The first iteration proved unusable due to an accidental 200% scaling error made during use of the software. The corrected coils fit on 6" square segments of circuit board. An outside vendor manufactured the coils on printed circuit boards; no further treatment was necessary.

RFID controller board. This board was designed with a socket to hold the matchbook-sized tag reader as if it were a single component. Matt Reynolds, a student with years of RF experience, recommended the use of triac-output optoisolators to switch the antenna signals. These permitted a low-voltage PIC processor to control the 140VAC signals without risk of damage or interference. An important constraint was to ensure that the antenna coil was not switched during the tag reader's communication with the coil; this would disrupt communications between the board and the tag, resulting in a faulty read or write operation. The PIC monitored a Clear-to-Send line on the tag reader, switching the coils as soon as this signal was detected. Upon detection of a valid tag, the PIC sent the ID to the adjacent LED driver board.

LED driver board. By this point, designing LED drivers had become second nature. This board used a PIC to control four 4x8 LED matrices. Every tag in the system was

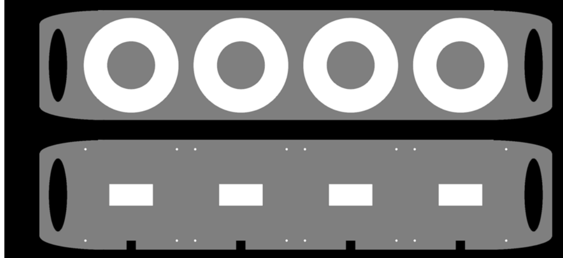

Figure 6.14: IDPad surface layout

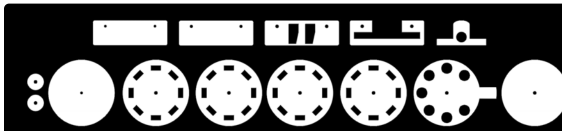

Figure 6.15: Laser cut layers for switch and revolver objects

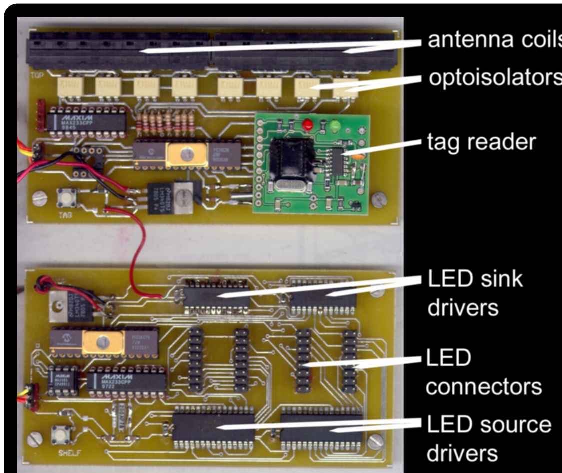

Figure 6.16: RFID controller board (top) and LED driver board (bottom)

associated with a unique graphical pattern. When a tag ID was received from the RFID controller, this tag was displayed; if no tag was present the display went blank.

Pad surface. As mentioned previously, the enclosure for this system was intended only to permit testing of the electronics prior to installation in a shelving system. Thus it was shaped to fit a row of four antenna PCB's, with the addition of handles at the ends for carrying. It was designed and cut one afternoon on the laser cutter, and was then assembled using standard machine screws.

Simple token objects. The objects used upon the surface of the IDPad were at first intended simply as handles for manipulation of the RFID tags. Because the wedge tags contained very directional antennae, they had to be held perpendicular to the field lines in order to be detected. A simple enclosure of clear acrylic was laser-cut; the name "bishop" reflects its similarity to a game piece. Another design which used several layers of clear acrylic was shaped like a loaf of bread.

Parameter control objects. During experimentation with the token objects, it was discovered that a strong magnetic field in the proximity of the wedge tags prevented communication with them. This was because the wedge tags, unlike the card and disc tags, used antennae with ferrite cores to concentrate field lines through their antennae. Magnetizing the ferrite material displaced the resonant frequency of the tag from 125kHz, thus preventing the tag's power circuit from operating.

This property inspired the design of objects with multiple physical states that could be distinguished wirelessly. These "parameter control objects" contained multiple tags and a number of small permanent magnets. The mechanisms inside the objects moved the magnets in relation to the tags, selectively disabling all but one of them. The "switch" object operates exactly like a toggle switch, enabling one of two tags embedded within. The "revolver" operates similarly, permitting selection of one of eight tags. One fortunate outcome of the magnet-suppression technique was that because the magnets were attracted to the individual tags, the switch and revolver objects tended to "snap" to discrete positions that represented different states.

Figure 6.17: Bishop and bread objects on IDPad surface

Figure 6.18: Objects on surface of IDPad

6.4 Discussion

Application concepts

Time did not permit development of an application that took advantage of IDPad's properties. However, for example's sake, consider an example inspired by MediaBlocks, Music Bottles, and other tools for media navigation and playback.

In a performance of electronic music, prerecorded audio samples are layered and modulated into dense compositions; they are accelerated, slowed, distorted in pitch and volume, echoed, appended, truncated, and concatenated. IDPad could allow all of these operations to be carried out through manipulation of the physical objects. Parameter control objects could serve to represent either objects or operations, while adjusting the controls could trigger switching between samples or could modulate effects performed on the samples. Because tag position is identified as well as identity, operations could also be distinguished by *location* rather than by object; thus common operations like recording and playback might be separated spatially on a control platform.

Also, the revolver object could be used to control relative rather than absolute identity. Thus, an digital album of 100 photographs could be browsed with an eight-position wheel by simply maintaining a position index in software. Another interesting aspect is that because the objects retain physical state, the modification of object attributes can occur without requiring the presence of the interactive surface.

Tangible interfaces vs. augmented physical interactions

Both SensePad and IDPad allow arbitrary information sources to be associated with physical objects; both have similar size, displays, and identification technology. Differences in the physical objects employed, however, illustrate a key distinction between this work and other "tangible interfaces."

The manipulation of collocated, untethered physical objects has been the focus of perhaps half of the tangible interface systems developed to date. The metaDesk, Triangles, Illuminating Light, Urp, Music Bottles, and the Media Blocks "sequencer rack" all respond to the presence and relative positioning of a set of token objects. These objects are designed to serve as easily manipulable and identifiable tokens for control of the interface, and are thus integral and inseparable components of the system. Any prior or concurrent physical uses (such as containment of liquids in Music Bottles) are incidental to their main function as interface control objects. The IDPad is another example of this type of system. The IDPad objects bear no resemblance to existing physical objects, and are useless without the system that recognizes them. By contrast, SensePad and TouchCounters are fundamentally based upon the existing uses of objects as physical containers to transport, organize, identify, and protect their contents. IDPad represents a new form of interaction, while SensePad represents an augmentation of existing interaction.

Category	Projects	Interactions	Object properties
Object manipulation systems	• Triangles • Bottles • I/O Bulb • metaDesk • MediaBlocks (sequencer) • IDPad	• Manipulation of collocated, fixed set of objects • Operations include selection, binding, sequencing, relative positioning	• Objects designed as part of system, lack functionality apart from use in the system • Sensing technology, extent of physical representation, and ergonomics
Distributed awareness systems	• InTouch • MediaBlocks (slots) • ambientRoom • Ambient Displays • musicBox • Network Meter	• Remote monitor-ing/transmission of data • Closely attended or passively monitored • Uni- or bidirectional communication	• Object represents either specific or generic data source • Object's properties must changes percep-tibly to indicate changes in data
Augmentation systems	• PPP • TouchCounters • SensePad	• Existing uses of object are maintained; digital information "augments" uses	• Tags and readers or recognition systems are fitted to existing objects

Table 6.1: Categorization of past Tangible Media Group projects

The success of these two types of systems depend on somewhat different factors. Compelling applications of IDPad, should they ever be developed, will hinge upon the digital functionality of the interactive system to which they are connected. The success of usage history displays like SensePad will depend upon the ability to provide and record useful data without negating the existing objects' functionality.

7 Evaluation and Future Work

This section describes the main contributions of the systems presented in the previous three chapters, and describes opportunities for future work.

7.1 FishFace and ShakePad

The FishFace and ShakePad projects demonstrated the feasibility of simple, self-contained interfaces with physically coincident input and output. They revealed the intrinsic attractiveness of bright, responsive graphics coupled closely to physical gesture. They also highlighted the sensitivity of the perceptual experience to slight variations in software implementation. The appeal of FishFace and ShakePad suggests that usage history displays can be interesting or even enjoyable to operate.

Because neither was applied to recording and displaying long-term histories of use, the immediate practical utility of FishFace and ShakePad remains unknown. Essentially, they are proofs-of-concept for both environmental and portable display devices.

To make FishFace and ShakePad better indicators of usage history, several technical improvements would be necessary. Increasing their "memory span" from seconds to hours or days would allow them to mediate interactions between temporally separated collaborators. To augment environments, longer-range sensing could be employed to detect body-scale motion. To augment portable objects, low-power displays such as "electronic ink" could be used to extend battery life.

Both FishFace and ShakePad are independent modules that can represent usage history without network connectivity. Reducing their physical size would increase the range of objects they could augment. In the long term, these modules could conceivably be miniaturized into independent, microscopic clusters that react to pressure, light, or temperature fluctuations. Such devices could be brushed onto surfaces like paint, instantly augmenting surfaces with history information.

7.2 TouchCounters

The TouchCounters system showed that tracking of containers rather than loose items confers several benefits. First, this approach simplified technical implementation, as numerous, heterogeneous, consumable items did not have to be tracked individually. Second, an existing physical categorization scheme was preserved and exploited. (Grouping items in a container is, in a sense, the ascription a common attribute to each item.) Finally, the array of evenly spaced containers was used to structure the visual presentation of multiple displays in parallel.

TouchCounters raised several questions regarding the complexity appropriate for a display of usage history. If usage history displays are simply visual representations of physical wear, are annotation or remote visualization operations necessary? Does passive tracking of objects constitute an invasion of privacy? Finally, what should be the balance of automatically triggered data presentation and the direct manipulation or control of data?

Controlled user testing of TouchCounters would be a strong step towards determining the appropriate complexity and feature set for a given environment. While the existing interaction modes were general enough for early testing, some customization to specific users, rooms, or industries would ultimately become necessary. Evaluation by real users might well result in reducing the number of features.

To conduct such long-term tests, the system would have to be made much more robust. Use of IDPad's contactless tracking system, rather than mechanical connectors, would greatly improve reliability. This technology would also allow the easy extension of the tracking system to a variety of object storage systems, including bookshelves, file folders, media storage racks, etc.

7.3 SensePad and IDPad

The SensePad project attempted to show that surfaces designed for storage and manipulation of objects could be redesigned as interfaces for information query and retrieval. In addition, it intended to demonstrate the low cost of augmenting objects through materials tagging technology. Although the tagging technology was not sufficiently developed for these goals to be fully realized, two elements of SensePad were significant steps towards lightweight interaction interfaces.

First, the contactless tags allowed identification and tracking of objects without swiping, scanning, or orientation towards a camera or other reader device. Second, the “digital shadows” were an effective mechanism for rapidly and unambiguously correlating attribute data with their physical referents. While the technique of collocating graphics with objects was used in the metaDesk, MediaBlocks and Luminous Room projects, only specific, custom-constructed objects were employed; SensePad showed this technique could be extended to augment objects with pre-existing functions and associations.

The IDPad prototype realized several of SensePad’s objectives with the notable exception of continuous position tracking. IDPad showed responsive, object-specific graphical underlays in action, and pioneered the use of rewriteable identification tags. While IDPad was intended only as a test platform for an augmented shelving system, it led to such a different direction that its original application was never implemented. The token objects with their parametric controls are better examples of tangible interfaces than displays of usage history.

In addition to debugging SensePad’s tagging and tracking technologies, greater effort could be spent to incorporate these technologies into the design of new types of furniture. New types of desks, tables, and shelves could be designed with integrated displays and tracking coils. This “computational furniture” would have immediate applications in retail environments, offices, libraries, and process automation, as well as in homes.

An opportunity not yet explored involves inverting the token/reference frame relationship of SensePad and IDPad. By making the tag reader into a compact, handheld object, the device could be used as a means to embed messages and annotations in objects throughout interior environments. Museums, kitchens, and bookstores could be annotated collaboratively by their users.

8 Conclusion

8.1 Summary

In this document, I have presented the concept of "situated displays of usage history," computational mechanisms for portraying the past history of physical objects and spaces. I have described the relationship of this concept to past and current research in human-computer interaction, and have outlined the space of devices that could potentially exist. From this space I have developed several specific prototypes that illustrate issues in design and application of these devices. I have also discussed social, economic, and psychological constraints upon the deployment of these systems.

From this work a few main points should be clear.

- Situated usage history displays are feasible and useful augmentations of existing environments.

- The tracking of human beings and of physical objects are parallel and complementary approaches, and many technologies are available for these purposes.

- To permit "lightweight interactions," situated usage history displays must be designed with particular sensitivity to preservation of existing patterns of use and understanding.

- In general, situated usage history displays should be used first to visualize data at the point of collection, rather than in a remote space. While more sophisticated functionality may be justifiable under some circumstances, the increased learning, orientation, and development time should be weighed carefully.

8.2 Personal vision

As one whose experience straddles research and product development, my tone throughout this document has been either analytical or pragmatic. The perspective of the artist—one compelled to express or elevate through creation—has been absent from my discussion. In closing, however, I will mention an aesthetic vision that has influenced my personal interest in this work.

I have a mental image of an entirely new sort of architecture or environmental space. This world is reactive, dynamic, and even organic, despite being man-made. In this world, all objects are alive and are responsive to human presence and action. Devices are omnipresent but unobtrusive, with individual visual indicators that can be seen only if examined at close range. As one steps away to survey a larger space, these tiny blips of light combine to form larger, flickering constellations.

One navigates this space as easily as one strolls through breezy, moonlit woods, gathering cues from brightness, color, spatial arrangement, and movement. Objects react instantly to human contact, just as frogs jump and croak when people approach. The environment continually evolves and breathes as a reflection of the people who inhabit it. “Sensors” and “displays” are symbiotic extensions of the structures that support them, and are found wrapped around pillars and walls like ivy vines. “Pixels” exist not upon planar, Cartesian surfaces, but speckled throughout space and buried deep within materials.

The behavior of each element of this world is extremely simple, but through the unpredictable interaction of its thousands of components, a world emerges that is magical, surprising, and alive. These crude prototypes are a far cry from approaching this vision, but perhaps they represent a small step.

A Appendix: Additional Interactive Devices

Tactile Pad and Network Meter Prototypes

This section presents two additional modular device prototypes that employed a variety of new sensing and display technologies. Conducted in an atmosphere of exploration and self-education, this work illustrates some elements of interaction relevant beyond the limited domain of the artifacts themselves. In addition, these projects served to develop technical proficiency that facilitated the implementation of later projects.

COMPUTATION	PIC microprocessors
DISPLAY	electromagnet arrays LED matrix displays
SENSING	electric field sensing acceleration sensing network data sensing
COMMUNICATIONS	serial interfaces radio communications

Table A.1: Interactive technologies explored in early prototypes

A.1 Tactile Pad

This project, the first I implemented at the Media Lab, was the construction of a haptic output mechanism for a colleague. While it has little direct relation to the other research, it is included for completeness and to highlight the evolution of the prototyping techniques throughout this work.

Background

Tom White, a student in the Aesthetics and Computation Group, had developed an input device using a silicone bladder filled with an opaque liquid (soy sauce). When a user pressed this bladder against the glass supporting it, a camera mounted underneath the glass could detect the shape of its "contact patch." This shape could then be used for the control of interactive graphics. Having developed an input device based on "liquid haptics," Tom's attention turned to the output side.

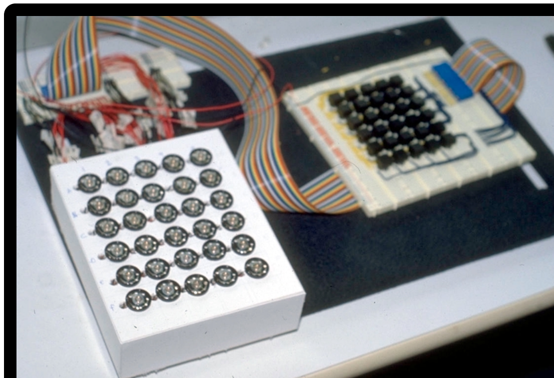

Figure A.1: "Lilliput" tactile array by Andrew Dahley

The idea of a high-resolution, responsive, tactile output array is often raised in discussions of haptics, but most employ mechanical pistons or vibrators. One example is the "Lilliput," an array of vibrating speakers controlled by a similar array of buttons.

Tom hoped to use magnetorheological fluid, which hardens during exposure to a magnetic field, as a kind of "haptic display." A square, flat bladder containing this material would rest atop an array of electromagnets. A computer could then control these electromagnets in order to generate patterns of stiff material throughout the bladder. The physical patterns would be the tactile equivalent of graphics. As a warm-up project, the author agreed to construct a prototype of this device.

Implementation

The system hardware was constructed in several days. Magnetorheological fluid was obtained from Lord Corporation, then a Media Lab sponsor; electromagnets were ordered by mail. The electromagnets were mounted in the grid holes of a sheet of pegboard bought at a hardware store. A total of 64 electromagnets were placed in an 8x8 grid. In order to maximize the field couplings through the material in the bladders, the electromagnets were wired in a "checkerboard" pattern, such that half of the alternate cells had reverse-polarity.

The electromagnets were driven by a set of transistors on a breadboard. These provided the electromagnets with high-current power, but could be controlled by low-power signals from the electronic circuitry. As each electromagnet had to be addressed independently, a separate driver transistor was used for each of the solenoids. The transistors were controlled by a set of eight octal latches that could "sample and hold" data from the PIC microcontroller. (The technique of "row/column addressing" was used in later projects to control LED arrays.)

Powering the array of electromagnets was not a trivial problem. Since each electromagnet drew 350mA current, the complete array could draw over 20 amps. A car battery was considered as a power source, but eventually a behemoth 24V, 20A power supply was obtained. This monstrosity, which lacked an enclosure and weighed over 30 pounds, posed the various hazards of fire, electrical outage, and crushing injuries.

An enclosure for the system was made by cutting and gluing pieces of black foam core around the circuitry; the bladder rested on the top surface. An LED display was installed for debugging, but was never fully implemented.

Figure A.2: Electromagnet array for Tactile Pad

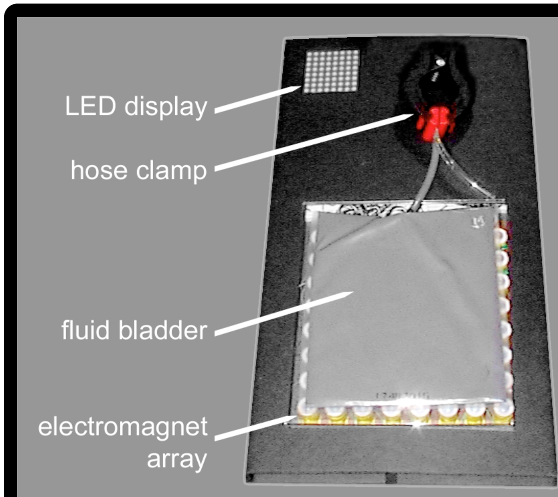

Figure A.3: Components of Tactile Pad

Performance

In designing this project, the author underestimated the difficulty of implementing a high-current drive system. Inductive loads on the solenoids caused a variety of electrical noise problems. Eventually, these problems were eliminated, but the system's performance remained unsatisfactory. The solenoids used were simply not strong enough to generate ridges of noticeable depth in the bladder fluid. However, dynamic changes were easier to detect; Tom White took advantage of this by creating "graphics" modes in which invisible forces swept across the pad array.

Discussion

Visual, auditory, or tactile displays can be used to communicate dynamic information. Visual displays, the most common, include monitors, LED displays, and projectors. Speakers, beepers, etc., are auditory displays, and tactile displays are usually motorized objects or vibrating actuators. (see table) The Tactile Pad was the only project to employ a non-visual output modality. Are haptic or vibrotactile actuators more appropriate for the task of distributed visualization?

The auditory transfer of information is usually serial. Since, in general, users must initiate contact with tactile displays to obtain information, touch is often serial as well. This is confirmed by Roger Whitehouse, designer of tactile maps for the blind:

> Sighted individuals can scan their surroundings and simultaneously become aware of the multiple possible structures and destinations within their purview; theirs is a complex and three-dimensional spatial understanding of the relationships of all of these things. The blind individual, on the other hand, experiences the environment sequentially, discovering first the door handle, then the door, then the frame, then the wall, then the wastebasket, then the chair, and so on. (Whitehouse99)

Certainly, tactile displays may be appropriate for surfaces that are not directly visible, such as furniture surfaces such as bench tops and seat backs. However, visual displays seem best for the representation of multiple, spatially distinct, moderate-bandwidth information streams. This is because engagement with such displays can be very rapid, information transfer can occur in parallel streams, and interference with physical manipulation tasks is low.

	Visual	Auditory	Tactile
spatialization of multiple sources	HIGH	MED	HIGH
physical contact required	NO	NO	YES
data bandwidth	HIGH	MED	LOW
interference with human-human communication	MED	HIGH	LOW
interference with concurrent manipulation tasks	MED	LOW	HIGH

Table A.2: Suitability of common output mechanisms for distributed visualization

A.2 Network Meter

Of the projects presented here, the Network Meter was both the most rapidly implemented and the longest in continuous operation. The device is a simple display that shows the level of network traffic between the Media Lab and the external Internet. It is a logarithmic graph that scrolls to the left as time passes.

Background

The Network Meter, which never had an official name, was installed at the suggestion of Jon Ferguson. Jon held a systems administration position at NeCSys, the Media Lab's Network Computing Systems office. To monitor the utilization of the lab wide Internet connection, Jon had written a Unix program that computed the current network bandwidth. Its output was an ASCII graph that advanced vertically up the screen of a workstation in the NeCSys office. Inspired to make this information available to the passersby, Jon inquired about the use of the LED displays he had seen in the TouchCounters project.

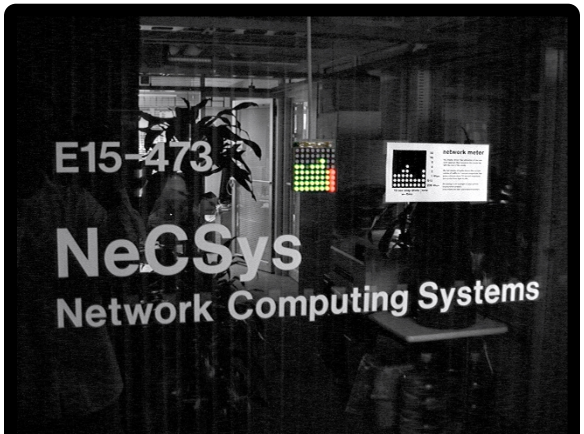

Figure A.4: Network meter in hallway window

Implementation

Adapting an existing label for this purpose was quite straightforward, and was done within a few days. The existing Unix script continued to perform the computation; Jon simply added a routine to send the display graphics commands via serial cable. PIC code was written to represent temporal data using a histogram. The vertical axis indicated the logarithmic bandwidth, which was normalized to a maximum value; the horizontal axis represented time. The rightmost column of the display showed the current bandwidth, and was updated every few seconds. Every 10 seconds, the average value during that time scrolled to the left by one column.

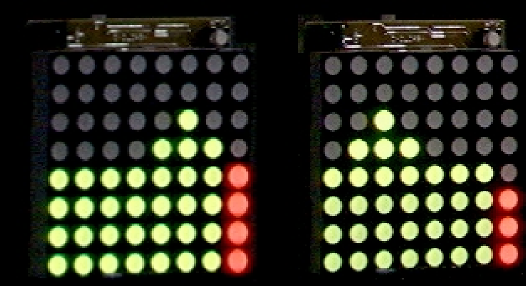
Figure A.5: Rightmost column displays current value; columns scroll left every 10 seconds

Performance

The Network Meter has been extremely robust in practice, and has been displaying data almost continuously since its installation. Systems whose electrical connections never change seem, in general, far easier to construct than reconfigurable systems.

Discussion

The electronic labels designed for the TouchCounters project had a modular design that made them easy to re-purpose for use in both ShakePad and Network Meter. This highlights the value of modularity, not only for products but for prototypes as well. Modularity in design implies both flexibility and scalability; FishFace could easily have been replicated to cover entire walls.

The Network Meter was hardly the first device to display network traffic in a public space. The Live Wire, created by artist Natalie Jeremijenko, was a ceiling-mounted device from which dangled a gyrating tube of rubber. For every Ethernet packet that passed over the local network, the tube twitched, thus providing an evocative physical indicator of network activity. Likewise, the Tangible Media Group's Pinwheels and Water Lamps are well-suited towards the physical manifestation of dynamic information streams.

The Network Meter, a simple visual readout, featured a far less radical representation than any of these projects. However, it did raise an important issue regarding the visualization of temporal data at multiple scales simultaneously. The network traffic value was refreshed every second or so, but were the display to advance every second, only a few seconds could be represented on the display. Thus it was necessary to divide the display area of the device into "current" and historical data.

If the display had been physically larger—a row of 100 LED matrices instead of one—this would not have been necessary. Such a structure would also have permitted the integration with auxiliary information, such as an event calendar.

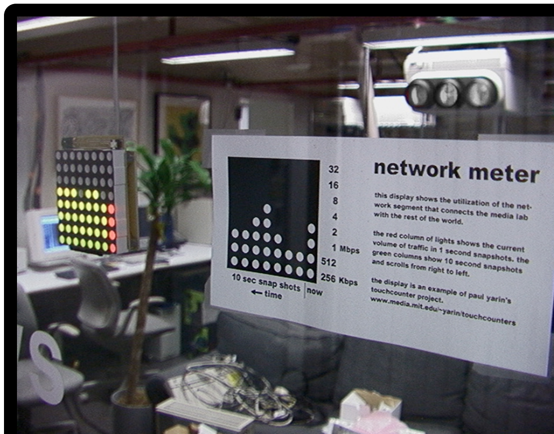

Figure A.6: Network meter legend

B Appendix: Project Influence Map

FishFace
FishFace was fairly original in its use of a bright, responsive, portable display. It borrowed the use of embedded PIC processors from Triangles.

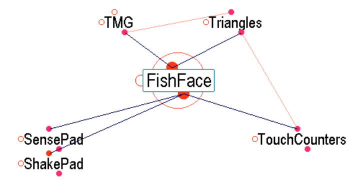

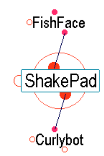

ShakePad
ShakePad was directly based upon FishFace, and can be seen as leading to Curlybot, a self-contained, proprioceptive device.

SensePad
SensePad utilized graphical underlays, employed in both metaDesk and MediaBlocks, as well as tagged passive objects as in MediaBlocks. Music Bottles also employs tagged bottles, and the IDPad tracks tagged objects with low-resolution graphical shadows as well.

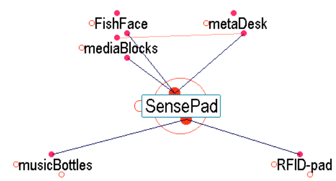

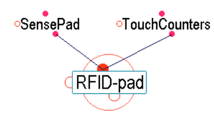

IDPad
IDPad employed tagged objects, as do SensePad and TouchCounters; it also shared the responsive visualization. New, however, are the parameter controls on the individual objects.

TouchCounters
TouchCounters drew heavily on past work: ambient fixtures for persistent, peripheral visualization, MediaBlocks for tagged objects, Triangles for physi-cal/digital connectors and embedded PIC's, and FishFace for responsive visualization. The Network Meter and IDPad are outgrowths, as is Strata's architectural visualization.

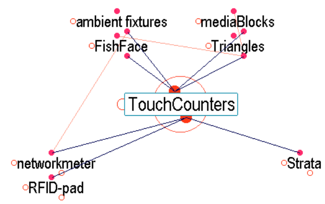

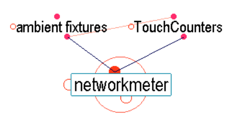

Network Meter
The Network Meter borrowed hardware directly from TouchCounters, but is based on an older concept from ambient fixtures: a single-purpose, statically-mapped display for continuous, remote data.

C Appendix: Design Parameters

Quantity of measure	
preservation of existing work practice	*existing actions vs. learned behavior*
usage type measured	*time-in-use vs. discrete events vs. value levels (temperature, movement, network traffic)*
criticality of information	*optional, augmentation information (boxes) vs. critical-path information (e.g. medication)*

Physical form	
physical distribution of displays	*scattered, object-centered vs. central cluster vs. environmental*
integration with other physical structures	*occasional attachment vs. removable label vs. integrated into design of object*
physical location of stored data	*local (FishFace/ShakePad) vs. buffered (SensePad) vs. network storage (TouchCounters)*

User interaction techniques	
degree of user involvement	*explicitly controlled (passive) vs. context-triggered (reactive) vs. automatic visualization*
modedness in interface	*persistent mapping (simple counting) vs. multi-modal (annotation+correlation functionality)*

Visualization	
temporal representation	*immediate (closely coupled) vs. historical (stored memory)*
responsiveness/dynamism	*dynamic, real-time interaction vs. slowly changing static display*
relativity of comparison	*single object in isolation vs. spanning multiple objects*
generality	*abstract or symbolic vs. literal or specific*

References

1. Brave, S. and Dahley, A. (1997) inTouch: A Medium for Haptic Interpersonal Communication. In *Extended Abstracts of CHI '97*. Atlanta: ACM Press, 1997.
2. Card, S., MacKinlay J.D., Shneiderman, B., eds. (1999) Information Visualization. In *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann, 1999, pp.1-38.
3. Cooper, Ken. (1998) Personal conversation, September 1, 1998.
4. Cramer, J. and Spilker, B., eds. (1991) *Patient Compliance in Medical Practice and Clinical Trials*. New York: Raven Press, 1991.
5. Dahley, A., Wisneski, C., and Ishii, H. (1998) Water Lamp and Pinwheels: Ambient Projection of Digital Information into Architectural Space. In *Extended Abstracts of CHI'98.* Los Angeles: ACM Press, 1998.
6. Dennett, D. (1998) *Brainchildren: Essays on Designing Minds.* Cambridge, MA: MIT Press, 1998.
7. Edwards, Paul N. (1994) From 'Impact' to Social Process: Computers in Society and Culture. In Sheila Jasanoff et al., eds., *Handbook of Science and Technology Studies*. Beverly Hills, CA: Sage Publications, 1994, pp. 257-285.
8. Feiner, S., MacIntyre, B., and Seligmann, D. (1993) Knowledge-based augmented reality. In *CACM* 36(7), July 1993, pp. 52-62.
9. Fishkin, K., Moran, T., and Harrison, B. (1998) Embodied User Interfaces: Towards Invisible User Interfaces. In *Proceedings of EHCI'98*. Heraklion, Crete, September 1998. In Press.
10. Fitzmaurice, G. (1993) Situated Information Spaces and Spatially Aware Palmtop Computers. In *CACM* 36(7), July 1993, pp.38-49.
11. Fitzmaurice, G., Ishii, H., and Buxton, W. (1995) Bricks: Laying the Foundations for Graspable User Interfaces. In *Proceedings of CHI'95.* ACM Press, 1995, pp. 442-449.
12. Fletcher, R. (1996) *Low-Cost Electromagnetic Tagging Technology for Identification, Sensing, and Tracking of Objects.* MS Thesis, MIT Media Lab, 1996.
13. Gorbet, M., Orth, M. and Ishii, H. (1998) Triangles: Tangible Interface for Manipulation and Exploration of Digital Information Topography. In *Proceedings of CHI '98*. Los Angeles: ACM Press, 1998, pp. 49-56.

14. Healey, J. and Picard, R. (1998) StartleCam: A Cybernetic Wearable Camera. In *Proceedings of the International Symposium on Wearable Computing*. Pittsburgh, Pennsylvania: ISWC, 1998, pp. 42-49.

15. Hearst, Marti A. (1995) TileBars: Visualization of Term Distribution Information in Full Text Information Access. In *Proceedings of CHI'95.* ACM Press, 1995, pp. 59-66.

16. Hill, W. C., Hollan, J. D., Wroblewski, D., and McCandless, T. (1992) Edit Wear and Read Wear. In *Proceedings of CHI'92.* New York City, New York: ACM Press, 1992, pp. 3-9.

17. Ishii, H. and Ullmer, B. (1997) Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In *Proceedings of CHI '97*, Atlanta: ACM Press, 1997, pp. 234-241.

18. Ishii, H., Wisneski, C., Brave, S., Dahley, A., Gorbet, M., Ullmer, B. and Yarin, P. (1998) ambientROOM: Integrating Ambient Media with Architectural Space (video). In *Summary of CHI '98*. Los Angeles: ACM Press, 1998, pp. 173-174.

19. Kirsh, D. (1995a). The Intelligent Use of Space. In *The Journal of Artificial Intelligence*, 73(1-2), 1995, pp. 31-68.

20. Kirsh, D. (1995b) Complementary Strategies: Why We Use Our Hands When We Think. In *Proceedings of 7th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum, 1995.

21. Minsky, M. (1985) *The Society of Mind.* New York: Simon and Schuster, 1985.

22. Norretranders, T. (1998) *The User Illusion: Cutting Consciousness Down to Size.* Viking Press, 1998.

23. Pashler, H. (1998) *The Psychology of Attention.* Cambridge: MIT Press, 1998, pp. 1-99.

24. Philips Hitag Information Page. (1999) Available at: http://www.semiconductors.com/identification/products/contactless/hitag

25. Picard, R.W. (1997) *Affective Computing.* Cambridge, MA: MIT Press, 1997.

26. Poor, R. (1999) The iRX 2.1 ...where atoms meet bits. Available at: http://www.media.mit.edu/~r/projects/picsem/irx2_1/

27. Raskin, J. (1999) Presenting Information. In *Information Design,* Robert Jacobson, ed. Cambridge, MA: MIT Press, 1999, pp.341-348.

28. Rekimoto, J, and Nagao, K. (1995) The world through the computer: computer augmented interaction with real world environments. In *Proceedings of UIST'95.* ACM Press, 1995, p. 29.

29. Resnick, M., et al. (1998) Digital Manipulatives: New Toys to Think With. In *Proceedings of CHI'98*, 1998, pp. 281-287.

30. Reeves, B. and Nass, C. (1996) *The Media Equation*. New York: CSLI Publications, 1996, pp. 219-226.

31. Smith, J. (1996) Field Mice: Extracting Hand Geometry from Electric Field Measurements. In *The IBM Systems Journal*, Volume 35, No. 3&4, 1996.

32. Streitz, N., Geißler, J.,and Holmer, T. (1998a) Roomware for Cooperative Buildings: Integrated Design of Architectural Spaces and Information Spaces. In *Proceedings of CoBuild'98: First International Workshop on Cooperative Buildings*. GMD, Darmstadt, Germany, 1998.

33. Streitz, N., Konomi, S,H.-J. Burkhardt (Eds.): (1998b) *Cooperative Buildings: Integrating Information, Organization, and Architecture.* First International Workshop, CoBuild'98, Darmstadt, Germany, February 1998. Available at: http://link.springer.de/link/service/series/0558/papers/1370/1370000v.pdf

34. Turk, M. (1997) *Workshop on Perceptual User Interfaces* (website). Available at: http://research.microsoft.com/PUIWorkshop97/

35. Ullmer, B., Ishii, H. and Glas, D. (1998) mediaBlocks: Physical Containers, Transports, and Controls for Online Media. In *Proceedings of SIGGRAPH '98*. Orlando, Florida: ACM Press, 1998, pp. 379-386.

36. Underkoffler, J., and Ishii, H. (1998) Illuminating Light: An Optical Design Tool with a Luminous-Tangible Interface. In *Proceedings of CHI'98.* ACM Press, 1998, pp. 542-549.

37. Want, R., Fishkin, K., Gujar, A. and Harrison, B. (1999) Bridging physical and virtual worlds with electronic tags. In *Proceedings of CHI'99*. ACM Press, 1999, pp. 370-377.

38. Weiser, M. (1991) The Computer for the 21st Century. In *Scientific American*, 265(3), pp. 94-104.

39. Weiser, M. (1993) Some Computer Science Issues in Ubiquitous Computing. In *Communications of the ACM*, Vol. 36, No. 7 (July 1993), Pages 75-84.

40. Weiser, M. and Brown, J. S. (1997) Designing Calm Technology. In *CHI'97 Workshop on Research Issues in Ubiquitous Computing*. Atlanta, GA: ACM Press, 1997.

41. White, T. (1998) *Introducing Liquid Haptics in High Bandwidth Human Computer Interfaces. MS Thesis*, MIT Media Lab, May 1998.

42. Whitehouse, R. (1999) The Uniqueness of Individual Perception. In Robert Jacobson, ed. *Information Design.* Cambridge, MA: MIT Press, 1999, pp. 103-129.

43. Wisneski, C., Ishii, H., Dahley, A., Gorbet, M., Brave, S., Ullmer, B., Yarin, P. (1998) Ambient Displays: Turning Architectural Space into an Interface Between People and Digital Information. In *Proceedings of CoBuild'98: First International Workshop on Cooperative Buildings*. GMD, Darmstadt, Germany, 1998.

44. Yarin, P. and Ishii, H., (1999) TouchCounters: Designing interactive electronic labels for physical containers. In *Proceedings of CHI '99*. Pittsburgh: ACM Press, 1999, pp. 362-369.

45. Zimmerman, M. (1989) The Nervous System in the Context of Information Theory. In R.F. Schmidt and G. Thews. eds. *Human Physiology*, 2d ed. Berlin: Springer-Verlag, 1989, p. 172. As quoted in (Norretranders98).