Folk Computing: Designing Technology to Support Face-to-Face Community Building

Richard Daniel Borovoy
Bachelor of Arts, Computer Science
Harvard University, June 1989

Master of Science, Media Arts and Sciences
Massachusetts Institute of Technology, September 1996

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN MEDIA ARTS AND SCIENCES
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2002

Signature of Author — Program in Media Arts and Sciences
January 11, 2002

Certified by — Mitchel Resnick
LEGO Papert Associate Professor of Learning Research
Massachusetts Institute of Technology

Accepted by — Andrew B. Lippman
Chair, Departmental Committee on Graduate Students
Program in Media Arts and Sciences

Folk Computing: Designing Technology to Support Face-to-Face Community Building

By Richard Daniel Borovoy

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on January 11, 2002 in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Media Arts and Sciences
At the Massachusetts Institute of Technology

ABSTRACT

Creating common ground in a community of people who do not all know each other is a chicken-and-egg problem: members do not share enough common ground to support the kinds of conversations that help build it. "Folk Computing" technology is designed to help build community in informal, face-to-face settings by giving users a playful way of revealing shared assumptions and interests. Drawing on the communicative process of folklore, Folk Computing devices facilitate the creation, circulation and tracking of new, digital forms of lore. These digital folklore objects serve as social probes: they circulate among people with whom they resonate, thereby revealing the boundaries of groups who share the underlying beliefs, knowledge and experiences that give the lore meaning.

Folk Computing uses technology to enhance the community building functions of folklore in three important ways: it supports the circulation of more interactive and media-rich lore, it reduces the social and cognitive costs of folklore creation and circulation, and it enables detailed visualizations of how pieces of lore circulate through a community. This thesis will explore the potential of Folk Computing through a design rationale for three new technologies, ranging from computationally augmented name tags used at conferences (the Thinking Tags and Meme Tags) to devices with which people can create, trade and track animations and simple games (the i-balls), used over several weeks by the population of a K-8 public school.

Thesis Supervisor: Mitchel Resnick
Title: LEGO Papert Associate Professor of Learning Research

Doctoral Dissertation Committee

Thesis Advisor

Mitchel Resnick
LEGO Papert Associate Professor of Learning Research
Massachusetts Institute of Technology

Thesis Reader

Justine Cassell
Associate Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Thesis Reader

Judith Donath
Assistant Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Thesis Reader

Henry Jenkins III
Ann Fetter Friedlaender Professor of Humanities
Massachusetts Institute of Technology

Acknowledgements

First and foremost, I thank my advisor, Mitchel Resnick. His vision drew me to the Media Lab, and his constant insight, support and good humor sustained me here. I also thank the other members of my thesis committee – Justine Cassell, Judith Donath and Henry Jenkins – who introduced me to many of the ideas in this thesis, and encouraged me to engage with those ideas in a manner that is both rigorous and playful.

I collaborated on the Folk Computing work with a large number of extremely talented people. In the many war stories that I tell about these projects, the heroes tend to be the same few people: Brian Silverman, Fred Martin and Tim Gorton. Also thanks to Sunil Vemuri, Chris Hancock, Jeff Klann, Matt Nototwidigdo, Matt Klicka, Brian Knepp, Chris Lyon, and Michelle McDonald for crucial contributions. Others, like Mike Rosenblatt, Tim McNerney, Alice Yang, Jan Malasek and Bill Thies, have had the generosity of spirit to jump in at the eleventh hour before a big event to complete some critical-yet-unfinished piece of the project.

Media Lab sponsor event-based projects like the Thinking Tags and Meme Tags put a huge burden on the administrative staff. Particular thanks to Carolyn Stoeber and Felice Gardner for their incredible effectiveness and professionalism.

I am very grateful for the support my work has received from Motorola as a Motorola Fellow, as well as from SEGA, LEGO, Intel, Tomy and HP.

Thanks to all the great people who have made being a part of the Epistemology and Learning/ Lifelong Kindergarten Group so enjoyable and stimulating, including David Williamson Shaffer, David Cavallo, Randy Pinkett, Vanessa Colella, Michelle Hlubinka, Dianne Willow, Robbin Chapman and Kwin Kramer.

I would like to thank my parents, Brenda and Roger, and my sister, Amy for giving me an early lesson in the value of lore as a means for creating a strong sense of "us."

Finally, thanks to my wife, Erin, whose shared understanding I value above all else.

Table of Contents

1 Introduction to Folk Computing 8

1.1 Folklore as a Medium for Social Probing 9

1.2 Limitations of Folklore as a Social Probe 12

1.3 Using Technology to Enhance Social Probes 14

1.3.1 Integrating Technology into Informal, Face-to-Face Conversation 14

1.3.2 Reducing Barriers to Participating in the Folklore Process 15

1.3.3 Supporting the Construction and Circulation of Resonant Texts 16

1.3.4 Visualizing the Circulation of Folklore 19

1.4 Background 20

1.4.1 Folklore 20

1.4.2 Other Models of Circulating Texts 22

1.4.3 Mutual Knowledge and Common Ground 23

1.4.4 Icebreakers 27

1.4.5 Technology to Support Face-to-Face Communication 28

1.4.6 Map of Primary Theoretical Elements 32

2 Using Tech to Support Talk: The Thinking Tags 33

2.1 Designing for the Creation of Mutual Knowledge 34

2.2 Design Technology to Augment Informal, Face-to-Face Conversation 36

2.2.1 Programming 36

2.2.2 Thinking Tag Display as Talk 37

2.2.3 IR Communication and Conversational Distance 40

2.3 Designing for Resonance 41

2.4 Events as Community Laboratories 42

2.4.1 User Community 43

2.4.2 Maintenance 43

2.4.3 Leverage 43

2.4.4 Technology Development 44

3 Creating a "Folk Culture" Culture: The Meme Tags 45

3.1 Technology and Activity Overview 46

3.2 Minimizing the Barriers to Participating in the Folklore Process 47

3.2.1 Reducing the Need for Immediate Authorship 47

3.2.2 Authorship 49

3.2.3 Choosing What Lore to Perform 50
3.2.4 Performance 53
3.2.5 Choosing What to Lore to Adopt 57
3.2.6 Remembering Lore 57
3.2.7 Sanctioning the Folklore Process 58
3.2.8 Results of Lowering the Barriers to Folklore Participation 58
3.3 Community Mirrors 60
3.3.1 Designing Technology to Track the Face-to-Face Circulation of Folklore 61
3.3.2 Designing the Community Mirrors 64
4 Designing an Environment for the Construction and Circulation of Resonant Objects: The I-Balls 70
4.1 Overview of Technology and Activity 70
4.2 Supporting the Construction of Resonant Objects 72
4.2.1 Animations 72
4.2.2 Wait Rules 74
4.2.3 Jump Blocks 77
4.2.4 Mutating Toward Resonance 79
4.3 Supporting the Oral Circulation of Digital Folk Objects 83
4.4 Building Visualizations of Patterns of Resonance 84
4.4.1 I-Ball Specific Visualizations 86
4.4.2 Visualizations that Reflect Individual I-Ball Exchanges 90
5 Folk Computing and Education: Supporting an Ecology of Learners 95
5.1 New opportunities for studying complex systems 96
5.1.1 I-Ball Integration with School Ecology 96
5.1.2 Constructing StarProbes to Investigate School Ecology 99
5.1.3 Examples of Students as Ecologists 100
5.2 New approaches to teaching traditional subjects 104
5.2.1 Spelling 104
5.2.2 Science 106
5.2.3 Data Analysis and Mathematics 107
5.2.4 Computational Ideas 109
6 Conclusion and Future Work 110
6.1 Designing an Ethnographic Study 110
6.2 Designing a Larger Folk Computing Trial 115
6.2.1 Phosphori 115

6.2.2 MemeMail 116
Bibliography 119

1 Introduction to Folk Computing

Glory glory hallelujah,
Teacher hit me with a ruler.
Hit her on the bean
With a rotten tangerine
And she ain't no teacher any more.

A surprisingly number of adults remember singing some version of this song, or others like it, during their childhood (McCabe 1998)[1]. Of course, such subversive songs were not sung at school assemblies, performed on television, or taught to kids by their parents. Instead, like all children's folklore, they were passed on from child to child via performances during recess, on the bus, and in the bathroom.

Research offers many theories about the function of such folklore: it is a way to assert one's independence, a safe outlet for anti-social impulses, and a means of exploring taboo topics (McCabe 1998). However, the role that folklore plays in creating a feeling of fellowship among those who perform it is probably most responsible for why lore like the "Battle Hymn" parody is memorable. After all, people do not just remember the words to this type of song; they often have vivid memories of whom they sang it with, and of the satisfying sense of community that the performance engendered. The joyful performance of the song affirmed to the singers and their audience that they shared some of its underlying beliefs – such as "school is oppressive" – which led to a feeling of fellowship and still greater enjoyment.

As children grow into adults, they may shift from singing song parodies to telling jokes or circulating "in" references, but their use of folklore to explore and celebrate the experiences, beliefs and interests they share with others remains the same. For example, after a popular Seinfeld episode was aired that used the phrase "yada yada yada" to mean "and so on and so forth", this figure of speech made its way into a large number of conversations.[2] Its use and recognition in an exchange was a way for people to confirm to each other not just that they Seinfeld fans, but also that they shared some of the Seinfeld ideology lying beneath the phrase –

---

[1] Chances are, the later you were born, the more violent the fate of the teacher in the song. When I sang this in the 1970s, the end of the verse had evolved into "I hid behind the door with a loaded '44 and the teacher is no more." In the current climate of school violence, one would probably be expelled if overheard singing those words.

[2] Although this phrase's origin was in the mass media, its usage in everyday conversation spread by word of mouth, and therefore it counts as folklore.

for example, a dismissive attitude toward a mundane level of reality considered unworthy of one's attention.

The previous examples show how folklore can affirm the shared nature of a set of beliefs among a group of people who are already aware of what they have in common. Folklore can also help groups of relative strangers uncover this common ground. Some of the same jokes, stories and in-references that confirm and celebrate what an pre-established group has in common can also allow people to probe for particular shared beliefs among others they do not know. For example, a punk rocker's hair or wardrobe may serve a kind of in reference to identify himself to others who share certain ideological commitments.

There is concern, however, that folklore, along with the sense of community it engenders, is an endangered species, due to a type of habit destruction. Folklore thrives when people interact face-to-face in informal, unstructured settings, but this type of unstructured social time is becoming a scarce resource in people's lives – particularly the lives of children (Hofferth and Sandberg 2001). Furthermore, folklore tends to flourish in social contexts that have some stability over time. Such continuity is less common in a mobile age where the family, the neighborhood, the company, the school and the country represent much more dynamic populations than they once did.

This thesis will present a series of technologies designed to support the construction, circulation, and visualization of novel digital forms of folklore in dynamic, informal face-to-face environments, in an effort to help users build a strong sense of community. The use of technology for this purpose it meant to be challengingly ironic, since technology is often seen as antithetical to all that is unstructured and intimate.

1.1 Folklore as a Medium for Social Probing

How do two individuals go from knowing nothing about each other to having an intense feeling of "shared understanding?" How does a group of relative strangers develop a sense of the important beliefs, values, experiences and knowledge that unite and divide them? Among strangers, direct communication about such deep issues is often considered taboo (Emerson 1969). However, without some exploration about these foundational assumptions, it is difficult to move toward a more knowing relationship.

Various indirect methods of communication have evolved in informal, face-to-face settings to help strangers overcome the barriers to friendship or collegiality while still maintaining appropriate social distance. As Irving Goffman, the foremost scholar in the area of face-to-face interaction, says:

> When two teams establish an official working consensus as a guarantee for safe social interaction, we may usually detect an unofficial line of communication which each team directs at the other. This unofficial communication may be carried on by innuendo, mimicked accents, well-placed jokes, significant pauses, veiled hints, purposeful kidding, expressive overtones, and many other sign practices. (Goffman 1959)

Researchers have particularly focused on the use of humor as a "self-disclosure and probing tool" (Kane, Suls et al. 1977) that people use to indirectly express their own underlying beliefs and to explore the beliefs of others:

> Persons are also naturally interested in fathoming what intentions, motives, and values... others possess. Standards of propriety may prohibit a person from directly asking others about these matters. A less direct approach would be to make a humorous remark that communicates the source's interest: presumably, if the target laughs and later reciprocates with a similar form of humour, the social relationship has moved toward more intimacy without committing either party in such a way that he or she could be called to account for their actions (Kane, Suls et al. 1977).

It turns out that folklore – "tradition-based communicative units informally exchanged in dynamic variation through space and time" (Toelken 1979)– can serve as a powerful medium for social probing. Most people have some experience with the concept of folklore. This may come from hearing an "urban legend" about the person who bought a Porsche for a dollar from the philandering owner's angry spouse (Brunvand 1981). They may remember the games they used to play as children – such as "Four Square," "Marco Polo," "Sardines," "One Potato, Two Potato," and the many variants of "Tag" and "Marbles" – that they learned from older peers and passed on to younger ones (Opie and Opie 1959). These experiences may leave one with the sense that folklore is something fun, but frivolous. Therefore, it is surprising for some to learn about the important role scholars believe folklore plays in helping to define and support the groups in which it circulates.

Folklore scholars have explored how humorous texts[3] that circulate orally can be a socially safe way for whole communities to establish a sense of what they have in common, and a sense of commonality. For example, McDowell explores children's riddling as a "cultural form given over to examination of the composition and boundaries of a culture." (McDowell 1979) Of

---

[3] Throughout this thesis we use "text" in the cultural studies sense to mean any activity, event or performance that is subject to interpretation

course, the domain of folklore is not limited to humor. Folklore researchers have examined how a wide variety of circulating texts – including legends, games, and even recipes for ethnic food – help establish community identity. (Dundes 1989)

Although there is considerable debate in the literature about what the different types of folklore have in common, for our purposes we will focus on their "resonant" quality. We define resonance as the visceral experience of a shared understanding between people that arises from a particular type of interaction. Resonance is the outcome of a transaction between a person who offers a text she finds personally meaningful and a person or group who demonstrably accepts that offer, confirming that they also find it meaningful. Members of the accepting group also experience "audience resonance" among themselves, as evidenced by the feeling of fellowship that can be shared by members of a stand-up comedy audience (Martineau 1972).

People engaged in resonant interactions often assume they find the text meaningful for the same reasons. Therefore, their shared understanding does not just include their belief that the text is meaningful, but the underlying beliefs, experiences and knowledge they use to make meaning of the text. In sum, a text is resonant when one person makes an assertion about its personal meaningfulness that another accepts, thereby establishing the shared nature of the assumptions and experiences necessary for making sense of it. The more that participants in the transaction identify with these underlying assumptions, the stronger the experience of resonance.

By focusing on the resonant aspects of folklore, we can think of it as a kind of socially appropriate "probing tool" whose circulation can establish the multiple, overlapping resonant groups that comprise communities. For example, the Cultural Studies scholar Henry Jenkins discusses this process in terms of the videotapes of favorite episodes that television fans share with each other. He suggests "fans have chosen these media products from the total range of available texts precisely because they seem to hold special potential as vehicles for expressing the fans' pre-existing social commitments and cultural interests." When a tape is exchanged between fans, the transaction between the tape provider and recipient establishes the shared nature of these commitments and interests, and gives the participants a positive experience of resonance. As Jenkins says, "what the videos articulate is what the fans have in common: their shared understandings, their mutual interests, their collective fantasies." (Jenkins 1992)

The "pre-existing" nature of fans' commitments and interests mentioned by Jenkins is one of the properties that gives folklore leverage as a means for establishing a shared understanding. The process does not require an elaborate shared experience to create a sense of shared meaning. Instead, it taps into personal meanings previously and independently established, and reveals them to be shared. Folklore draws further leverage from the fact that the beliefs and interests underlying it are established as shared without ever having to be directly articulated.

1.2 Limitations of Folklore as a Social Probe

It is legitimate to talk about resonant groups as being both created and revealed. Although the group exists before as a set of people who have some beliefs in common, it is not aware of its existence. This sudden awareness of the shared nature of these beliefs is the revelation brought on by a resonant text, which then leads to the creation of a group's self-concept. There are some limitations of folklore as a probing tool that reveals and reifies like-minded groups, however. First of all, it is difficult to discover the set of people within a population that resonate with a particular piece of folklore. Consider humorous lore. There are ample cues – for example, laughter – for a person to determine whether a joke he has told resonates with someone who is listening. However, the sample of data he collects about who resonates with the joke is likely to be biased by the make-up of his friends and acquaintances. As these people pass this joke on to others, the previous teller will likely get no feedback about which people or even how many people found it funny.

In fact, there is evidence to suggest that people have poor understandings about the circulation patterns of folklore. Consider a few examples drawn from email folklore. Recently, the Boston Globe reported on some students there who posted a large banner outside their dorm with the words "All your base are belong to us", a somewhat nonsensical phrase from a videogame that had miraculously become a very popular piece of Internet folklore. But the Harvard students, who had heard about the phrase from a friend over email, said they didn't realize its popularity: "The scary thing is that we thought it was our private joke" (Denison 2001). Similarly, a Wired Online article reported that the person who created a parody of Florida's infamous butterfly ballot was surprised when he found out the email he sent to 30 friends had quickly spread to 130,000 people. He said "I had no idea that my email had circulated around the world" (Dean 2000).

In making sense of the circulation patterns of their own folklore, children often believe that a friend invented a folk game they recently learned. They are often surprised and incredulous when told it was passed down from Roman times, for example, as was the case with Marbles (Knapp and Knapp 1976).

Finally, consider an example that is not about folklore per se, but documents people's impoverished intuitions about circulation within their social network. In Stanley Milgram's famous research, subjects were asked to guess the number of intermediate acquaintances it would take to relay a letter from a random person in the United States to another random person whose name and address was known. People guessed 100 on average, while the actual number was closer to six (Milgram 1967).

Another limitation of folklore is the finite expressiveness of texts that can circulate in informal, face-to-face settings. One must be careful here, because texts like the Iliad and the Odyssey originally circulated orally, and a gifted storyteller can create a vibrant world through words and gestures alone. However, the ubiquity of PowerPoint slides in academic presentations attests to many people's desires to augment their oral exchanges with rich media. Fan communities have pushed this even further, particularly at fan conferences, where "the creation, exhibition and exchange of videos" plays "a central role in solidifying and maintaining the fan community." (Jenkins 1992). The problem is that most media technology has not been designed for use in informal, face-to-face conversational settings. In the case of video, unwieldy technology is still required to create, view and exchange it. Though some early adopters are willing to put up with these limits, most people are not.

While it is true that humor allows people to take "interpersonal initiatives that would otherwise be too risky" (Kane, Suls et al. 1977), humor and folklore do not eliminate these risks. In fact, the creation and exchange of folklore is still fraught with risks and social costs. One such risk, particularly in adult communities, is the appearance of frivolity. Because folklore is not explicit about the underlying beliefs it probes, many people do not understand the important role it plays in creating community. Therefore, participating in the folklore process can be judged insubstantial and unproductive – for example, consider many people's attitude toward email folklore. Since folklore is a kind of play – it allows people to try on and try out a variety of assumptions encoded in different folk texts without fully committing themselves – one can read adult reluctance to participate in folklore exchange in terms of the larger issue surrounding the appropriateness of play for adults.

The social probing afforded by folklore has an associated cost, which is particularly noticeable when someone introduces multiple probes in search of some resonance. For example, consider when two strangers meet and search for someone they know in common. Although this is not a folkloric exchange, it has similar qualities: as with a piece of folklore, a common acquaintance is used as an indication of a deeper and more meaningful set of commonalities. As most people know from experience, however, there is a point in the "Do you know so-and-so" game where the potential benefit of finding some commonality is outweighed by the cost in time and social awkwardness of cycling through multiple, unfamiliar names (at this point, the game is often terminated by one person saying "That name sounds familiar").

The final limitation of folklore we discuss here is the relatively high "cognitive load" required to participate in the folkloric process (Sweller 1994). The cognitive acts of creating (or modifying) a folk text, choosing a folk text to perform for a particular audience, performing it, and attending to and remembering the performances of others are all fairly demanding. If someone feels particularly weak at any one of these skills – for example, someone who does not remember jokes or feels he does not tell them well – then he will be unable to participate in the entire folklore cycle.

In the final analysis, there are enough risks and barriers to participating in the folklore process that many "folks" simply choose not to, thereby limiting the community building benefits it affords. These benefits are further curtailed by the inherent limitations of folklore discussed above, including the limits on what texts can be circulated in informal, face-to-face settings, and on how well folklore can be tracked as it circulates.

1.3 Using Technology to Enhance Social Probes

This thesis will explore Folk Computing technology: technology to help people create and affirm a sense of commonality and disparity between and within groups in their community by facilitating the construction, circulation, and tracking of resonant texts via informal, face-to-face conversation. Using the folkloric process as a point of departure, we will show how technology can be used to both support this process and overcome some of its limitations. One major focus will be on how technology can be integrated into informal, face-to-face conversation. We will also explore the role technology can play in: reducing social and cognitive barriers to the folklore process; supporting the construction and circulation of highly resonant texts; and establishing the circulation patterns of folklore texts. These themes will be introduced below using the three major Folk Computing devices as examples.

1.3.1 Integrating Technology into Informal, Face-to-Face Conversation

Since the purpose of Folk Computing technology was to support community building in informal, face-to-face settings, it was essential that it be well adapted for this environment, and even protective of it. Therefore, we designed the technology to work within the structures and processes of face-to-face conversation. As an example, consider the first Folk Computing device that we built, the Thinking Tags (see Figure 1-1).

The Thinking Tags were nametag-like devices designed to give participants at a gathering a simple measure of how much they had in common by comparing their answers to five multiple-choice questions. Questions were chosen to be relevant to the gathering. For example, at the Media Lab event where we did the first user trial, one of the questions was "How would you like to spend your fifteen minutes of fame? a) An appearance on Oprah; b) An interview on the front page of the New York Times; or c) Your home page linked to the main page of Yahoo." After programming their tags, when two people met face-to-face, the five LEDs on their tags showed one green light for every question they answered similarly, and one red light for every one they answered differently.

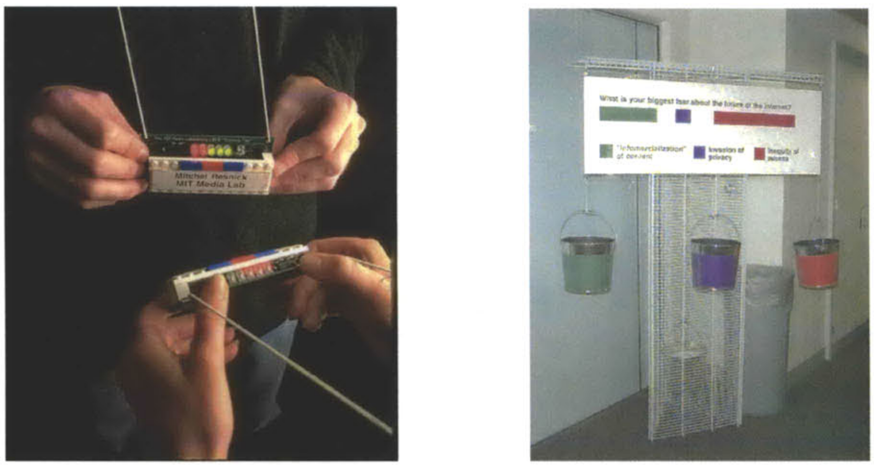

Figure 1-1 Thinking Tags (Left) and Bucket Kiosks (Right)

There were several ways we worked to insure the Thinking Tags helped conversation instead of hindering it. One of the large challenges was finding a social way for people to program their tags with their opinions. The easy solution of having people program their tags at some kind of PC kiosk seemed at odds with the surrounding social context. Instead, we created a system where people could program their Tags by dunking them in a paint bucket that represented their answer, hung from a large sign that stated the question. In this way, several people could program their Tags simultaneously while already beginning to discuss the questions.

1.3.2 Reducing Barriers to Participating in the Folklore Process

Folk Computing technology was designed to overcome some of the cognitive and social barriers to participating in the folklore process. For example, the second Folk Computing device we designed – called "The Meme Tags" – significantly reduced the costs and risks associated with creating and circulating folklore.

Participants could program their Meme Tags with short 64 character ideas, or "memes", that they believed in (e.g., "`If brute force isn't working, you're not using enough of it`"). When two people met, each person's tag displayed a meme that he or she subscribed to and that the other person had not yet seen (see Figure 1-2). If people liked the meme they saw on their conversation partner's tag, they could click a button and a copy of it would "jump" to their tag. We also created large screen "community mirrors" that showed visualizations of how memes were moving through the community in real time.

Figure 1-2 The Meme Tags (Left) and a Community Mirror (Right)

Meme Tags made it easier to participate in the folklore process by choosing which piece of lore to perform, by managing the performance, and by recording it for a recipient who wanted to pass it on to someone else. The dedicated folklore "channel" they created also made it easier and less risky for people to introduce a piece of lore, without having to worry about framing it appropriately; the tags themselves provided the frame, and, to some extent, sanctioned the activity.

1.3.3 Supporting the Construction and Circulation of Resonant Texts

Folk Computing devices supported the construction of folklore objects that had the potential for greater resonance than purely verbal lore, but could still circulate and replicate in informal, face-to-face contexts. For example, the third Folk Computing device we built allowed users to create animations and simple games that could then be passed from person to person. Figure 1-3 shows the graphical programming tool participants used to create their "i-balls" (short for information balls), which then ran on the key-chain sized i-socket.

Figure 1-3 The I-Ball Authoring Tool (Left), the I-Socket (Middle), and an I-Ball Exchange (Right)

An i-ball's ability to display animations and to exhibit a variety of interactive behaviors gave them more expressive power. For example, Figure 1-4 shows an i-ball authored by the school librarian that referred to Van Morison's popular song "Brown Eyed Girl". The right-most block of the program strip (top of figure) is an animation block that will display the text "my brown eyed girl" with two eyes that alternately wink at the viewer (bottom of figure). What is particularly interesting is that the two "wait rule" blocks to the left of this animation block ensure it will only play if the i-socket it is running on is owned by a brown-eyed girl. Otherwise, a more generic animation plays that says "van the man." The details about how this operated will be explored in Chapter 4. Here, we simply call attention to the way she was able to use the programming environment to explicitly specify an "in group" for the i-ball within the i-ball text itself. Computationally amplifying the ability of a piece of folklore to specify an in-group may have made the "1970" i-ball resonate even more strongly with this group. Of the seven people who got the "1970" i-ball from its brown-eyed female creator, four of them were brown-eyed females, and one of them passed it on to two more brown-eyed females.

Increasing the resonant potential of folklore texts while still maintaining their ability to circulate in face-to-face, informal settings is a challenge. In many ways, the i-balls were similar to more tangible forms of folklore, such as the kinds of paper airplanes and "cootie catchers" that kids make. Such objects have two limitations that the i-balls overcome, however: cost of replication and requisite expertise. Tangible folklore doesn't circulate as freely in informal settings because, unlike its purely verbal counterpart, it is not as easy to replicate and pass on. One of the great virtues of bits over atoms is this ease of replication. Kids could make an arbitrarily complex i-ball and then give away as many copies as they wanted out in the playground, without worrying about finding the right materials or tools.

The expertise required to make more complex folk objects creates another barrier to their easy construction and circulation. Bronner points out that "because the maker of objects is often a specialist, a gulf can exist between the maker and the viewer, who may not share or understand the skills of the craftsman." (Bronner 1986) This gulf limits the development of a true, decentralized community where everyone can participate in creating a space of resonant objects in which they can locate themselves in relation to each other. Folk Computing technology bridges this gap in several ways. First, computational design tools can scaffold the construction of complex objects, making a process available to novices that was once only accessible to experts. For example, the i-ball authoring environment makes it very easy to create an i-ball that consists of a single picture. Most people figured out how to do this in a few seconds time. From there, we designed the tool to make it a single step to turn that picture into animation. After that, it was relatively easy to turn that animation into a full computer program that exhibited a variety of behaviors.

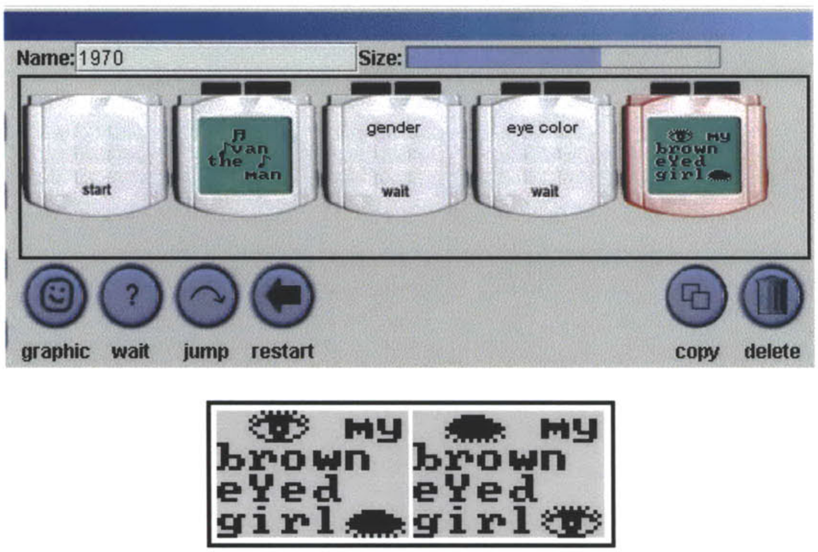

Figure 1-4 "1970" I-Ball Program (Top) with Animation Frame Detail (Bottom)

Folk Computing further reduced the gap between expert and novice builder by allowing recipients of any i-ball to "open it up", see how it works, and then modify it to suit their own needs. In fact, the author of the "1970" i-ball discussed above created it by modifying the program of another i-ball she received. This was possible because computational media made it relatively straightforward for us to preserve the i-ball program "genotype" for each i-ball "phenotype" that circulated. When someone was interested in understanding or changing the behavior of the phenotype, they could then access and modify the underlying program genotype. Not all media make it possible to go from phenotype back to genotype.

1.3.4 Visualizing the Circulation of Folklore

Specially designed digital folk objects that circulate through a computational medium are much easier to trace than traditional folklore that circulate via word of mouth, allowing people to see the entire population among which a piece digital lore has circulated. This is in contrast to the traditional egocentric – specifically "ego network-centric[4]" – view of a piece of folklore, where someone is aware of who they got it from, who they gave it to, and perhaps a few other exchanges they observed or heard about. In the case of the i-balls, participants were able to view different types of visualizations that showed how an i-ball spread from its creator through the user population. Figure 1-5 shows the two visualizations that we used to determine that the "1970 / Brown Eyed Girl" i-ball went mostly to brown eyed females. The top visualization is colorized by gender, where the blue nodes are females who got the i-ball, the purple nodes are males, and the links run from top to bottom connecting the person who gave the i-ball to the people who got it. The bottom visualization is colorized by eye color, where the brown nodes are people with brown eyes (one can immediately see that everyone who got the i-ball had brown eyes).

The ability to visualize the audience for a piece of folklore has important implications for establishing a resonant group. As explained previously, a group can have an experience of resonance when its individual members observe each other resonating with a particular text that has been offered. Visualizations allow community members to observe which other members resonated with a particular text when it was presented to them. Of course, this is missing the simultaneity of a group of people finding a new text meaningful at the same time as realizing that others find it meaningful. In the case of the visualizations, one usually finds out who else found it meaningful after she has found it to be meaningful herself. Nevertheless, well-designed visualizations appear to be capable of engendering a form of "audience" resonance in those who view them.

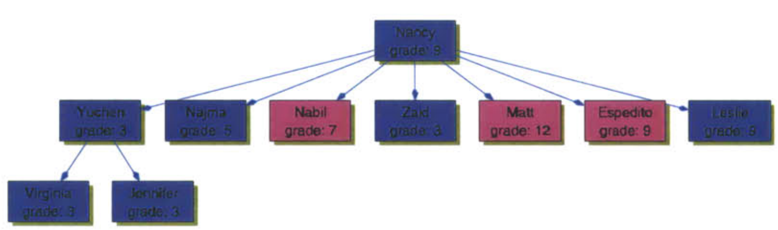

---

[4] An ego network is a social network analysis graph that shows a single person's relationships (Wasserman and Faust 1994)

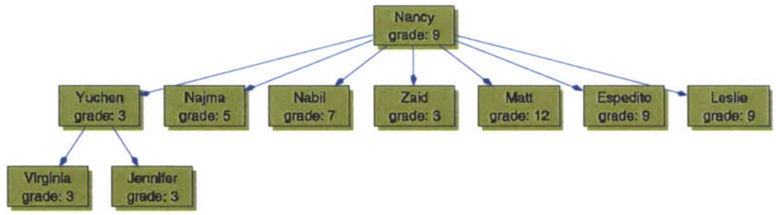

Figure 1-5 "1970" I-Ball Diffusion Visualizations Colored by Gender (Top) and Eye Color (Bottom)

1.4 Background

The following sections explore the various areas of research on which "Folk Computing" draws. The final section provides a succinct "map" of how the different theoretical elements fit together.

1.4.1 Folklore

Unlike most researchers, we were interested in folklore from a design perspective. Inspired by the role folklore plays in building community, we wanted to help people create, exchange, and track their own computationally enhanced folklore-like objects. Folklore scholarship is a contentious and heterogeneous field, however, where there is considerable disagreement even about the very definition of folklore. Therefore, based on a review of relevant research, we have formulated the following definition of folklore in a manner that is useful for our own design purposes. *Folklore is a self-organizing social system that helps groups of people reveal, experience, and extend their commonalities and connections via the circulation of adaptive, resonant texts.* The following sections elaborate elements of this definition and their origins in the literature.

1.4.1.1 Resonant Texts

We define resonance as the experience of a shared understanding that results when a group of people realizes that they all identify with the presuppositions of a particular text. Although this concept runs through the literature on how folklore helps constitute a community, it is never formally named. For example, Brunvand suggests "the legends we tell, as with any folklore, reflect many of the hopes, fears and anxieties of our time." (Brunvand 1981) Similarly, speaking about particular television shows that fans have used as the basis of their folklore, Jenkins says "there is already some degree of compatibility between the ideological construction of the text and the ideological commitments of the fans." (Jenkins 1992) Toelken discusses folklore as being highly "connotative," where connotation is the "attitudes, value judgments and implications" associated with the folk text. Furthermore, Toelken highlights the role connotation plays in creating "a shared sense of 'we.'" (Toelken 1979)

1.4.1.2 Self-Organizing Social System

A piece of folklore circulates widely because a large number of individuals find it meaningful enough to pass along. Unlike mass media, a small group of people does not determine what a large audience will receive. Instead, the folk group regulates its own consumption, and manages its own informal system of distribution. Oring summarizes the views of most contemporary folklorists by saying "[folklore is not] from the elite and their centers of political, cultural and commercial power, or from institutions of media communication...folklore cannot be legislated, scripted, published, packaged, or marketed and still be folklore."(Oring 1986)

Because the variants of a piece of folklore that resonate with large numbers of people will thrive, folklore can be seen as an evolving, adaptive system that puts selection pressure on the quality of resonance. Again, speaking about the fan community, Jenkins says, "[folk] songs are constantly being rewritten, parodied, and amended in order to better facilitate the cultural interests of the fan community." (Jenkins 1992)

1.4.1.3 Revealing and Experiencing Commonalities

Patterns of folklore circulation delineate groups of people with common deep-seated beliefs, experiences, fears, etc., and make knowledge of these groups available to the groups themselves. For example, Toelken discusses the outfits worn by most loggers in the American Northwest. While some of the garments are chosen for safety and utility, much of the uniform "constitutes a folk costume by means of which loggers belong and recognize each other."(Toelken 1979) One of Jenkins informants explains how group laughter in response to the performance of a piece of folklore reveals a pattern of shared resonance: "If you 'get' the joke, punch line or reference and laugh when the rest of the fan audience laughs, it reinforces the sense of belonging, of 'family', or shared culture." (Jenkins 1992) Of course, we all use this process when we share stories with people we do not know very well, and use their reaction as a way to gauge our commonality. Such practices do not just reveal some commonality; they can also provide a strong feeling of commonality. For example, Oring discusses how ethnic food, considered to be a type of folklore, can be used "to create a sense of community within groups, as well as to define and delimit boundaries between groups" (Oring 1986).

Folklorists use the term "folk group" to describe groups of people who are bound together by folklore. Jay Mechling defines "folk groups" as "face-to-face human groups wherein people use stylized communication to create the sense of a shared, meaningful world" (Mechling 1986). The "stylized communication" he refers to is folklore. Note that Mechling considers that the folklore is what gives rise to the sense of a "shared, meaningful world," not the other way around. Other researchers also emphasize the reciprocally reinforcing nature of folklore and the folk group. As Toelken says:

> One of the key features of a folk group will always be the extent to which its own dynamics continue to inform and educate its members and stabilize the group. Because the members share so much information and attitude, folk groups are what Edward T. Hall would call High Context Groups... whose members all see themselves as parts of a single community that 'knows'. (Toelken 1979)

Due to the resonant, non-explicit nature of folklore, however, patterns of shared belief can be revealed without the need for direct apprehension of those beliefs. In this way, groups of people who may not be well enough established, comfortable enough, or introspective enough to discuss these things directly can use folklore as a way to both uncover these commonalities and even to stand in for them in group discourse.

It is worth noting that because folklore works to both maintain old groups as well as establish new ones around newly discovered common interests, it contradicts a purely structuralist view that only allows for replication of traditional groupings. For example, the structuralist Mary Douglas – like many folklorists – considers jokes to reflect underlying social structures. However, unlike folklore researchers, she does not leave open the possibility that reflecting these structures back to group of people can change how the group sees itself, and ultimately change those same underlying social structures (Douglas 1999).

1.4.1.4 Extending Commonalities and Connections

Beyond the revelation of existing commonalities, folklore can help people who already know each other expand their shared commitments, as well as help people who do not know each other begin to find common ground. For example, McDowell discusses how telling riddles, another accepted form of folklore, " allows the child to expand his communicative network beyond the immediate circles of kin and close friends. New relationships are facilitated by the availability of riddling as a technique of communication between nonintimates." (McDowell 1979).

1.4.2 Other Models of Circulating Texts

1.4.2.1 Memetics

The ideas behind the nascent, quasi-academic study of "memetics" have something in common with the folkloric process. Richard Dawkins first introduced the term "meme" to suggest how ideas can spread and evolve through Darwinian selection (Dawkins 1989). He conceived of them as a kind of cultural gene. Memetic researchers have drawn parallels between folklore and memetics (Lynch 1996). However, this thesis draws on folklore theory rather than memetics for several reasons. First of all, the role of resonance is not as well defined in memetics as it is in folklore. Although one can imagine employing the concept of resonance in how someone decides whether to adopt a new meme, there is very little agreement in the memetic community as to what an appropriate memetic "fitness function" is. Also, the community-building implications of memetics have not been explored. For example, memetics offers no equivalent to the useful concept of the folk group.

If memetic theory was not useful for Tag design, why did we call them Meme Tags? We believed the concept of a meme was already familiar to many people, and would help them understand the activity. Furthermore, although the term meme was invented to sound like gene, it also sounds like "name", as in nametag.

1.4.2.2 Social Life of Documents

In their widely circulated 1996 article on "The Social Life of Documents," John Sealy Brown and Paul Duguid discuss the role that circulating documents play in constituting community:

> "People with shared interests use communications technologies (both hi- and low-tech) to help form themselves into self-created and self-organizing groups. To a significant degree, these are held together by documents circulating among members, each keeping each conscious of being a member and aware what others are up to." (Brown and Duguid 1996)

Brown and Duguid do not make reference to the folklore literature, although their points about the community-building and -defining role of written texts were very similar to the points that folklore researchers had made previously about oral texts. Their focus on written documents leads to useful explorations of the role of appendages like routing slips in helping people view a text's pattern of circulation. However, the written nature of documents also make this work less applicable to thinking about community building in face-to-face settings. For example, in Brown and Duguid's discussion, documents lack the connotative quality that is essential to folklore's role in establishing a sense of community around ideas and beliefs that are either too inchoate or too risky to state directly. Also, documents are not treated as evolving, dynamic entities that continue to search out new community boundaries, as they are in the folklore literature.

1.4.3 Mutual Knowledge and Common Ground

One of the purposes of Folk Computing technology is to help strangers in informal, face-to-face settings get some information about what they have in common. Everyone intuitively knows the importance of this task from countless conversations with strangers that start with two people trying to determine whether they share any significant people, places or things in common. It turns out that this precursor to more meaningful conversation has its roots in a basic theory of communication – namely, that all communication depends on an understanding of one's audience. As psychologists Robert Krauss and Susan Fussell state: "For people to communicate effectively… they must develop some idea of what their communication partners know and don't know in order to formulate what they have to say to them." Krauss and Fussell define this challenge as "the mutual knowledge problem", where mutual knowledge is "knowledge that the communicating parties share and know that they share" (Krauss and Fussell 1990). The condition that the two parties not only know, but also "know that they know," is essential to allowing interlocutors to successfully interpret their partner's utterances. Consider the example offered by Herbert Clark, one of the major formulators of the key role of mutual knowledge in language use, where Ann asks Bob "Have you ever seen the movie showing at the Roxy tonight?" (Clark 1992) To formulate this sentence, Ann must know what Bob knows about what is playing at the Roxy. To interpret it, Bob must know what Ann knows about what Bob knows what is playing at the Roxy. Clark has formulated the more inclusive concept called "common ground" -- defined as the sum of two people's "mutual, common, or joint knowledge, beliefs, and suppositions" – and calls it "a sine qua non for everything we do with others." (Clark 1996)

By designing technology to help establish common ground, we wanted to help transform strangers at a conference into potential collaborators. This is in keeping with Clark's claim that "aquaintedness comes in degrees defined largely by the type and amount of personal common ground two people have," where strangers have none, acquaintances have a limited amount, and friends have extensive common ground (Clark 1996). However, the normal process by which people uncover and create common ground can be cumbersome and time consuming. The high cost of establishing common ground led Krauss and Fussell, in their chapter on "Mutual Knowledge and Communicative Effectiveness", to provocatively ask – but not answer – "are there ways in which technology can reduce the difficulty of formulating what is mutually known?" (Krauss and Fussell 1990)

The circulation visualizations produced by Folk Computing technology help solve a variant of the mutual knowledge problem. Krauss and Fussell discuss several processes for establishing mutual knowledge, one of which is shared experience. In this way, "everybody thought the joke was funny" becomes mutual knowledge to an audience whose members simultaneously laugh at a joke. What happens when two people in different settings laugh at the same joke? Although they to some extent share a common experience, they do not know that they share it. Therefore, mutual knowledge has not yet been created.

Folk Computing visualizations help build mutual knowledge among a group of people who resonated with a common piece of folklore given to them by different people at different times, enabling them to experience a sense of "audience resonance" with each other. We have named this type of mutual-knowledge-building display a "μ-cue," where μ is the Greek letter "Mu", the first syllable in mutual. A μ-cue is just a label for something that already existed, however. Consider the example of a "4 Way" stop sign.

When two cars pull up to an intersection posted with "4 Way" stop signs (see Figure 1-6), their drivers immediately share a set of assumptions that that allow them to make appropriate sense of each other's actions. For example, from Driver A's perspective:

> Driver A knows Driver B has a stop sign. Therefore, if Driver B drives through the intersection without stopping, Driver A can confidently interpret this as an unlawful act, and respond appropriately (e.g. by honking, gesturing, etc.).
>
> Driver A knows Driver B knows Driver A has a stop sign. Therefore, if Driver A is inclined to ignore her stop sign, she can be confident Driver B will interpret it as an unlawful act, and she will be able to interpret his response accordingly.

Figure 1-6 Four Way Stop Sign as μ-Cue

In the above example, the two drivers have formed an ad-hoc "interpretive community" – the literary theorist Stanley Fish's term for a community that shares enough of the same knowledge, beliefs, and practices to be able to negotiate the meaning of a particular text (Fish 1980). Ordinarily, interpretive communities are thought to grow slowly and incrementally over long periods of time. However, the small "4 Way" sign-within-a-sign plays an important role in creating a "quantum leap" toward an interpretive community: by establishing the fact that "everyone has a stop sign" as common ground for the drivers, much of the meaning of intersection activity can then be shared.

The "4 Way" sign works by transforming an existing pattern of distributed knowledge into mutual knowledge, where distributed knowledge is knowledge that is shared, but not known to be shared (my definition), and mutual knowledge is "knowledge that is shared, and known to be shared." (Krauss and Fussell 1990) Specifically, the "4 Way" sign transforms a situation where everyone at an intersection is thinking "I must stop" – a pattern of distributed knowledge – to a situation where everyone knows everyone is thinking "I must stop", and everyone knows everyone knows everyone is thinking "I must stop" – mutual knowledge.[5] This last clause, while sounding ridiculous, is crucial. If we stop at "everyone knows everyone is thinking 'I must stop'", then each person might still think that while he or she knows that everyone has to stop, the other people do not know it (for more of this type of analysis, see (Clark and Carlson 1982)).

A μ-cue is an artifact that augments the edges of an interpretive community by transforming an existing pattern of distributed knowledge into mutual knowledge. Edges of interpretive communities occur both when a newly forming community struggles to establish some common ground (such as when four drivers pull up to an intersection), and when a newcomer attempts to acquire enough of the mutual knowledge of an established community to begin to make appropriate sense of its discourse. We define μ-cues in terms of community edges because in the center of these communities, members have already established a large body of mutual knowledge; it is on the outskirts that initiates need help in their search for common ground.

As another example of a μ-cue, consider the "To:" and "Cc:" fields at the top of an email message. These fields establish the message contents as mutual knowledge among its recipients. The diabolical "bcc:" establishes a very complex structure of mutual knowledge by saying "I want you to know. I want you to know they know. I don't want them to know you know. I want you to know they don't know you know."[6] Users take these μ-cues for granted, but they are a very powerful feature: one can imagine early designers of email systems seeing no need to provide information about who else has seen a message. These cues are powerful because they establish shared knowledge that users can draw on to craft their own messages and to interpret the messages of others. (Brown and Duguid 1996)

Perhaps the first instantiation of a high-tech μ-cue to build common ground between people in a face-to-face setting was for a research project on social uses of Personal Digital Assistants (Borovoy, 1993, unpublished). An application called "Face2Face" utilized information from two

---

[5] In this example, we assume that the drivers have a mutual knowledge of the meaning of a four-way stop sign.

[6] Many people fail to pick up on this last implication, and send a reply to the named message recipients, effectively "outing" the original sender for using a "bcc."

PDA users' address books to reveal whom they knew in common.[7] Face2Face was a play on the "Do you know..." game that two people often play when they first meet. It used computation to compare two people's social contacts more quickly than they could "by hand" and therefore offered a way to reduce the social cost of finding this commonality. The availability of the address book data was also key to making this interaction inexpensive.

1.4.4 Icebreakers

A Folk Computing activity has some things in common with a traditional "icebreaker", defined in Webster's as synonymous with a "mixer" – a "game, stunt, or dance used at a get-together to give members of the group an opportunity to meet one another in a friendly and informal atmosphere" (Webster 2001). There are some important differences, however.

As an example, consider the instructions for an icebreaking activity called "Name Tag Match Maker", listed on one of the many websites that feature such content (see http://www.businessfundamentals.com/IceBreakers.htm#Name Tag). The activity bears some resemblance to the Thinking Tag activity. Participants are instructed to put their names on a 5" x 7" card, and then write some different things about themselves in the four corners of the card (e.g. "in the upper left corner, write four things that you like to do"). The instructions then say:

> "When everyone finishes, have them mingle with the group for a few minutes. Without talking, they are to read the upper left corner of the other group members' cards. When time is up, they are to find one or two people who are most like them and visit for a few minutes. When time is up, they are to mingle again reading the upper right corner of the other group members' cards... To make sure everyone visits with several people, you could implement a rule that no two people can be in the same group more than once."

The most important difference between the "Name Tag Matchmaker" game and the Thinking Tag activity is that the former takes place *in vitro*, and the latter *in vivo*. The "Matchmaker" activity is set in an artificial environment with its own arbitrary time limits and rules about whom you can to talk. In contrast, the Thinking Tag activity runs in the background of a "living" social gathering, and is designed to work within its existing structures. Many people express discomfort with icebreakers like the Matchmaker activity because of their artificialness: by imposing so much arbitrary structure on a gathering, the activities compromise the "friendly and informal" atmosphere they are meant to create. The intruding structures of traditional icebreaking activities

---

[7] Of course, this application only revealed which people in their address books had the exact same first and last name, which meant there could be many false positives and negatives.

cannot easily be changed, however. The Matchmaker activity requires everyone to be doing the same thing at the same time not because this is socially desirable, but because it is necessary. The simultaneous execution of specific rules makes the awkward act of comparing one's tastes to those of a stranger more acceptable and more manageable.

The Thinking Tags take over the laborious and awkward elements of the communication protocol described in the instructions for the Match Maker activity. The tags handle the issues of synchronization and comparison. This frees up the participants to immediately make use of this information in conversations that fit in to their own conventions and interests.

1.4.5 Technology to Support Face-to-Face Communication

Even in the era of the Internet and the cell-phone, few would argue about the privileged role that face-to-face conversation has in the production and sharing of meaning. As the anthropologist Marjorie Harness Goodwin says:

> "Were an ethologist from Mars to take a preliminary look at the dominant animal on this planet, he would be immediately struck by how much of its behavior within a rather extraordinary array of situations and settings (from camps in the tropical rain forest to meetings in Manhattan skyscrapers), was organized through face-to-face interaction with other members of its species." (Goodwin 1990)

However, at the time of the initial Thinking Tag experiment in 1995, there were no examples of technology that supported face-to-face communication at informal social gatherings (in fact, there are still very few examples). Why hasn't there been more work on technology to support informal, face-to-face gatherings?

One reason may be that face-to-face communication is often considered the gold standard against which computer-mediated communication is measured (Hollan and Stornetta 1992). This type of research is focused on using technology to reproduce the physical world's rich set of social cues in the online world and has little to say about what technology can do when those cues already exist. However, having praised the virtues of face-to-face settings, some researchers have overlooked their limitations. Goodwin claims that for a variety of different reasons, social scientists have also neglected face-to-face conversation as a significant area of study (Goodwin 1990). Perhaps this has also contributed to the lack of attention technologists have paid to this area.

Another reason why there had been so few technological inroads into face-to-face communication might be a deep-seated belief that technology is antithetical to intimacy. Consider the current debate about people's use of cell phones in public places. An article called "Cell Hell" in the online Salon Magazine bemoans a proposed use of cell phones on airplanes, "eliminating one of the last oases of unconnected time." (Mieszkowski and Quistgaard 2000) The cell phone is just the latest example of a technology that is derided for encroaching on our humanity. However, it may appear especially threatening because, where technologies like television and the Web had to lure us to them, the cell phone was one of the first that was able to follow us into the restaurant, the park, the car – into our daily, personal world. With the Thinking Tags, we wished to prove that a piece of technology worn by every member of a large social gathering could enhance the feeling of community at a social gathering, not damage it.

The following two sections highlight relevant research into Folk Computing in the areas of online and co-present community building technologies.

1.4.5.1 Online Community Building

Most tools for online community building have a very different set of requirements than face-to-face tools. While face-to-face tools can leverage traditional communication modalities such as voice and gesture, online tools must carry the entire substance of the communication themselves. Early versions of such tools were limited in terms of the "dimensions" of face-to-face communication they represented. For example, on-line text-based discussion tools carried the text of a conversation, but couldn't reproduce much of the nuance carried by vocal inflection. Early video conferencing systems represented more visual and auditory cues, but failed to preserve subtler details like gaze awareness. In the last several years, some researchers have been working to reintroduce a variety of face-to-face social cues into tools for remote communication (Ishii, 1993 #41 and Viegas and Donath 1999).

Some researchers have been exploring interfaces for on-line "social visualization" (Donath 1995) that, like Folk Computing, try to give participants a window into larger community dynamics. Warren Sack's work on Very Large Scale Conversations (Sack 2000) provides members of a Usenet on-line community with a social network visualization showing who is responding to whom in the conversation. Judith Donath, et al's, research on "Virtual Fashion" (see http://www.media.mit.edu/~dc/research/fashion/) focuses on "tracking, analyzing, and visualizing cultural dispersion on the World Wide Web". Obviously, these approaches differ from Folk Computing in that they are designed to support on-line, not face-to-face, community building. They also differ significantly with regard to their fundamental unit of analysis. Folk Computing uses the transmission event – where two people purposefully exchange a text – as the basis for many of its social visualizations. This means that the computer doesn't have to do any complex parsing of text or human activity to determine where a connection between two people has occurred, which makes the visualizations more feasible and reliable. From the perspective of a community's comprehension of itself, that also means that these visualizations are based on events that are readily discernable and graspable by individuals. The experience people have creating their own Folk Computing lore (memes, i-balls, etc.) and exchanging it with others proves very helpful in making sense of the resulting complex visualizations of these exchanges.

1.4.5.2 Same Time, Same Place Technology

Although it is in the minority, research exists on using technology to support face-to-face communication and community. Some of this work focuses on "meeting support," where the goal is to help groups of people in a meeting to brainstorm ideas and/or make decisions. This is different from Folk Computing along several key dimensions. First, meeting-support technology usually plays a much larger role in mediating the dominant conversation. For example, the purpose of some tools is to ensure that people can contribute anonymously (Nunamaker, Dennis et al. 1991). Other tools, such as digital whiteboards, help users create novel representations of their ideas to make a conversation more productive (Stefik, Foster et al. 1988). Compared to these, Folk Computing devices work much more in the background of a community. Like folklore, they represent a type of stylized discourse not designed to carry the main thread of a discussion. They are also designed to work over a larger piece of space-time than the typical meeting, and with a larger population of people.

Another type of same-time, same-place technology research focuses on building awareness of other people's activity within a space. Technology exists to help people in a building track the whereabouts of others (Want, Hopper et al. 1992), and to get a sense of the activity they are engaged in -- e.g. Joe and Mary are standing around the water cooler, Chris is in her office reading email (Dourish and Bly 1992). Supporting this type of "background" awareness has something in common with the goals of Folk Computing, but the emphasis is very different. Rather than focusing on geography or activity, we have tried to make people aware of the ideas and beliefs that under-gird their communities, bringing some people together while simultaneously pushing others away.

COMRIS, or "Co-habited Mixed Reality Information Spaces" is a wearable technology research project that, like some of the Folk Computing research, focuses on augmenting conference-type gatherings (Velde 1997). The design metaphor they use is a parrot that sits on one's shoulder and whispers useful pieces of information into one's ear about what is going on in one's surroundings. Unlike Folk Computing, the main emphasis in COMRIS is on supporting individuals, not relationships or community. This is partly clear from the types of augmentation it provides, such as "agenda management." What makes it especially true is COMRIS' reliance on the "traditional" wearable computing approach, where the display is only accessible to the wearer (Starner, Mann et al. 1996). While Folk Computing attempts to help build common ground, COMRIS undermines it by making participants wonder whether a person they are talking to is listening to them, or to some private message about them relayed by the parrot.

There was a famous demo of "Cinematrix" technology at SIGGRAPH 98, where an auditorium full of people holding special paddles could collectively control a computer game projected on a large screen at the front (see www.cinematrix.com). However, this was a structured activity, with an audience in their seats giving their full attention to the technology.

With its interest in Things That Think, the Media Lab has inspired several artifacts that enhance face-to-face communication and community. The Galvactivator, created by Jocelyn Scheirer and Rosalind Picard, is the most similar to Folk Computing technology (see www.media.mit.edu/galvactivator). Like the Thinking Tags, it is a wearable device that gives people in a conversation some provocative information about each other – a measure of their arousal. Unlike the tags, however, the Galvactivator provides information that is about its wearer – not directly about the relationship between the wearer and the viewer. When Galvactivators are given to a large audience, they can also provide a kind of Community Mirror, like the Meme Tags. Again, what the mirror reflects is considerably different from the tags. One way to characterize this difference is in terms of time-scale. Arousal changes moment by moment. The common ground that an auditorium of flashing Galvactivators establishes is fleeting, and is perhaps potent because it is fleeting. The common ground that Folk Computing devices establish is based on resonance with people's deepest hopes, fears, and beliefs – things that endure over time – and it is potent because of this longevity.

Justine Cassell and her group have developed technologies designed to augment face-to-face socializing at a "literary salon". Some of this work helps create a sense of community by inviting people to join in playful coordinated activity (Cassell, Smith et al. 1999). Other parts of the work focus on novel means of message passing between participants in the shared space. These are different approaches to community building from Folk Computing.

Rob Poor created another round of augmented nametags for a Media Lab sponsor event. These functioned quite differently from the Folk Computing tags in that they were not geared toward augmenting human interaction. Instead, they provided a way to "personalize" technology to respond to the unique attributes of individuals.

1.4.6 Map of Primary Theoretical Elements

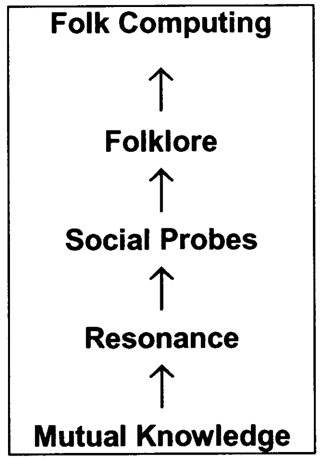

Figure 1-7 Theoretical Linkages

Although Folk Computing draws on a variety of theoretical elements, each one extends the power of the previous ones in a fairly systematic way, as shown in Figure 1-7. The most basic construct is mutual knowledge: Folk Computing technology aims to help affirm and grow the mutual knowledge of a group, as a way of creating a sense of community. Folk Computing helps establish mutual knowledge by supporting transactions involving resonant texts. When a text that represents an underlying set of knowledge, interests and assumptions is offered up and accepted as meaningful, participants in the transaction can assume that its presuppositions are mutual knowledge. Because these resonant texts allow people to explore the shared nature of these underlying assumptions without explicitly stating them, they serve as effective social probes for testing and affirming the beliefs of others in a group. Finally, folklore provides a medium in which these social probes can circulate in an informal, decentralized fashion, allowing all members to create social probes to explore the underlying beliefs and interests of a complex, dynamic community.

2 Using Tech to Support Talk: The Thinking Tags

Our objective with the Thinking Tags was to design a technology to help people in informal, face-to-face settings probe each other's beliefs and find common ground in a socially safe manner. The Tags were also designed to be a proof-of-concept that technology could play a positive and supportive role in such settings. At the time of their debut, there were no other examples of this.

Here is a quick review of how the Thinking Tags functioned: by allowing participants to "program" their badges with their multiple-choice answers to five opinion questions, two users could see a simple measure of how much they had in common by noting the number of flashing lights on their badges as they conversed. Each green light corresponded to a question they answered the same way. Each red light signified a question they answered differently.

The concepts and technologies of Folk Computing coevolved, and the Thinking Tags were created early on in this process. Therefore, these tags implemented what we now perceive as a scaled-down version of the full set of Folk Computing features. The following bullets compare these Thinking Tag features to the later technologies:

- **Small, fixed set of folk texts**: For each of five questions determined in advance, participants chose the answer with which they most strongly resonated from the three provided. With the Meme Tags and I-balls, there was a much larger set of potentially resonant texts that users could add to during the activity.
- **One-way "Star network" pattern of circulation:** With the Thinking Tags, texts did not circulate from person to person as they did with later technology. Instead, they moved from a central source – the bucket kiosks – to each of the participants. This is closer to a publishing sense of circulation, such as "newspaper circulation", than to decentralized, peer-to-peer sense of circulation in the domain of folklore.
- **Pair-wise μ-cues**: While later work on the Meme Tags and I-balls focused on establishing community-wide patterns of shared belief, the Thinking Tags focused on providing this information for two people in a conversation. Because the Thinking Tags were not responsible for exchanging texts between them (see above), their sole purpose was to insert a μ-cue into the communication. With the later technologies, the devices became the locus for the performance and exchange of

texts, and the μ-cues were moved to other locations. Therefore, the Thinking Tags were the only Folk Computing technology to explore the use of μ-cues directly in conversation.

- **Simple μ-cues:** The μ-cue used with the Thinking Tags was very basic: a display of the number of texts in the union and disjunction between each person's set of resonant texts. The Tags displayed information about the amount of common ground, but did not reveal any particular piece of common ground. In later work, visualizations got much more specific and elaborate.

The following two sections provide a design rationale for the Thinking Tags. A third section explores how our work with the Thinking Tags helped us discover the significance of resonant texts and their salient characteristics, which led to new objectives for later technologies. The final section discusses what we learned about the use of large social gatherings as a laboratory for exploring face-to-face community-building technology.

2.1 Designing for the Creation of Mutual Knowledge

In order to generate a μ-cue that reveals a pattern of knowledge distribution, the technology must be able to establish the distribution of knowledge across a set of users' devices. With Folk Computing technology, a piece of knowledge consists of whether or not a user subscribes to a particular text (and perhaps some information about who the text came from), and a μ-cue might show to whom a particular text has circulated. The cost of creating this requisite distribution of knowledge must be low. Ideally, μ-cues can be generated from knowledge that has already been collected and maintained for other purposes, as was the case for the original PDA application that used people's address books to determine who they knew in common.

For the Thinking Tags, there was no useful, preexisting knowledge base about participants to draw on for generating μ-cues. For the initial trial, we had access to the mailing addresses of the two hundred participants, but the promise of revealing to two strangers "You're both from New Jersey" seemed less than compelling. There was the possibility of trying to collect some data in advance in the form of a questionnaire, but we doubted many attendees would be willing to contribute information this way, especially in advance of experiencing the activity.

We settled on the idea of asking a small set of multiple-choice opinion questions to use as the basis for comparison. Such information would be easy to collect at the event, easy to compare, and relatively easy to display once it was transformed into a μ-cue. This strategy was in line with Stanley Fish's specific insights about the role of shared opinion in creating interpretive communities (Fish 1980). Fish suggested that revealing relevant common beliefs between two people might help them better understand each other. The questions in the Media Lab Thinking Tag event were chosen carefully after debate and community involvement to ensure not only their relevance but also their evocativeness. If the guests did not have strong feelings about any of the questions, then the tags would not really be uncovering patterns of preexisting beliefs, and the μ-cues would lack meaning.

By producing μ-cues, Thinking Tags are designed to not only displaying knowledge about the individuals in a conversation, but also knowledge about their *relationship*. The reason this is more powerful is similar to why Tom Erickson explains that the World Wide Web is powerful: it enables someone to get important information about someone else without having to "accrue a social debt to them" (Erickson 1996). Ordinarily, there is a substantial cost in terms of time and effort to establishing some meaningful common ground with someone through traditional approaches, such as awkward "Where are you from? Do you know..." conversations. This cost reflects the "chicken and egg" nature of the "Mutual Knowledge Problem" (Krauss and Fussell 1990) – it is hard to find it when you don't have some already, and you don't have any unless you find it. By reducing this cost, we hoped the tags might help people have more meaningful conversations, both in terms of quantity and quality.

There is a potential problem with computationally augmented nametags supporting mutual knowledge. Ordinary nametags establish mutual knowledge because the wearer knows what is on his or her own tag, and knows that the viewer will also know it. When the content of the tag is being generated dynamically, however, the wearer no longer knows exactly what information he or she is displaying. This could result in the uncomfortable situation where a viewer reacts to the contents of a wearer's tag, and the wearer does not know how to make sense of the reaction. For the Thinking Tags, we solved this problem by designing an augmentation scheme that could show the same contents on both guests' tags. That way, two guests conversing could look at each other's tags, and immediately know the contents of their own.

One might ask whether we really needed technology to do this activity. Could the same mutual knowledge be created if people simply wrote their answers to their opinion questions on their nametags and then looked at the tags of others to discover what they had in common? In addition to the cognitive overhead associated with doing this comparison, the problem here is that a guest would have no way of immediately knowing whether his or her conversation partner had done the comparison, even if the guest witnessed the partner looking at the guest's tag. While "physical co-presence" can establish the contents of a nametag as mutual knowledge (Krauss and

Fussell 1990), it cannot similarly establish the result of an analysis of the contents. In order to make this mutual knowledge, the analysis must be part of the contents.

2.2 Design Technology to Augment Informal, Face-to-Face Conversation

2.2.1 Programming

The key question: was there a way for people to program their answers into the tags that would preserve the informality and intimacy of the social gathering? The most obvious solution was to provide computer terminals with pointing devices where guests could pick their answers to the questions, and then have the answers transmitted to their badges via infrared. However, that would require people to take themselves out of their social environment so they could focus on the programming task. In a sense, the programming task would become a black hole, where the person doing it could not see out, and no one else could see in.

We came up with a novel solution to the programming interface problem when we considered it more carefully. Putting a traditional computer in a social environment creates two opposing spaces: the space contained by the room and the space contained by the screen. These spaces are on totally different scales and have totally different methods for interaction. One moves through the space of the room with one's feet, but through the space of the screen with one's hand. Many people can simultaneously participate in exploring the physical space, but only one can navigate the on-screen space at a time. In order to participate in the on-screen space, one must cease to participate in the social space.

Figure 2-1 The Bucket Kiosks

The solution came when we imagined what the programming interface would look like if we took it off the screen and spread it throughout real space. What if we took the five questions, each with three check-box answers, and distributed them through a room, so the programming space and the social space were intertwined?

Fusing programming and social space yielded the approach of the "Bucket Kiosks" (see Figure 2-1). Participants programmed their tags by walking up to each of five question kiosks and dunking their tags into the one of three buckets that corresponded to their chosen answer. The check boxes that would have divided a screen into meaningful areas were translated into buckets that divided a three-dimensional room into meaningful volumes of space. Instead of removing people from the social context, the bucket kiosks created a new sort of meaningful social interaction: guests engaged in conversations around the bucket kiosks, taught newcomers how to use them, and debated the questions as they made their selections. Bringing the programming task into the social space meant that the act of programming one's tag became simultaneously viable as a technology and visible to the other participants. It allowed the social and computational activities to reciprocally influence each other, and both activities became more meaningful.

2.2.2 Thinking Tag Display as Talk

We designed the Thinking Tags to function as a natural part of talk, producing a kind of utterance heard by two people beginning a conversation. Goodwin highlights several qualities of utterances in talk that we used as guides in our design, such as:

- Utterances are interpreted within a set of larger frameworks
- Utterances are crafted for a particular listener
- Utterances reflect "participant analysis of prior talk"

Our use of these concepts in Thinking Tag design will be discussed in the following sections.

2.2.2.1 Designing for Interpretation in Context

An important element of the mastery of natural language is "the ability to understand more than is explicitly said within a strip of talk by situating it within both indigenous frameworks of commonsense knowledge and the practical circumstances and particular activities in which parties to the talk are engaged" (Goodwin 1990). We wanted to ensure that participants could situate the contents of the Thinking Tag displays in their own activities and knowledge.

The challenge was to display a small amount of useful information in a social context without disrupting that context. A useful metaphor in our design process was the post-it note. The power of a post-it note does not come from the small amount of text one can scrawl within its boundaries. Rather, the power lies in the ability to temporarily affix a small amount of text in a context that gives it meaning for the period of time that the augmentation is useful. The key rules of the post-it note are: the content must be relevant to the context, the content must be bound to the context; and the content must not obscure too much of the context. Our challenge was to adapt the simple structure of a post-it note to a more dynamic environment.

The Thinking Tags needed to be able to show how many of the five questions two people answered in common. Initially, we considered a simple display consisting of one flashing LED, where the frequency of the flashing would correlate to how much two people had in common. This did not prove to be very visually stimulating, however, so we moved toward a display with multiple lights.

The most obvious idea was to map each question onto a single LED, the "correlation" display. If two people agreed on a particular question number, the light in that same position would glow green; if they disagreed, it glowed red. In Figure 2-2 Left, two people agreed on questions 2 and 5, and disagreed on the others. Therefore, lights 2 and 5 show up green, with the rest showing up red.

Among all the choices we considered, the correlation display provided the most information: theoretically, two people could directly perceive which issues they agreed and disagreed on, and could immediately launch into a discussion about them. In practice, however, this proved difficult for test subjects to do. In order to make sense of the display, users had to not only remember the five different questions, but they had to remember them *in order.*

The "summation" display algorithm provided a possible solution. In this algorithm, the question agreement data is reduced to a simple metric of how much the people have in common: each agreement is given a value of positive one, each disagreement is given a negative one, and the results are added together. If the result is positive, that number of lights is displayed in green. If it is negative, the number is displayed in red. If the result is zero, a single amber light is displayed. As indicated in Figure 2-2 Middle, two agreements and three disagreements sum to one disagreement, which is displayed as one red light.

As with the correlation display, test subjects had a hard time interpreting the summation display: the map between question agreement and display lighting pattern proved non-intuitive. In both these cases, instead of augmenting people's conversations, the tags took them hostage, demanding that the conversants spend their time trying to decode the display. This was a case of the post-it obscuring the underlying context.

In the "color sorting" display algorithm, the green and red lights (indicating answer matching and mismatching, respectively) are grouped together and displayed from right to left (see Figure 2-2 Right). This display method provided enough information to stimulate conversation and not so much as to stifle it. In fact, there was such a close match between people's expectations (including our own) about the display and how it actually behaved that the color sorting approach began to feel like the natural and obvious choice.

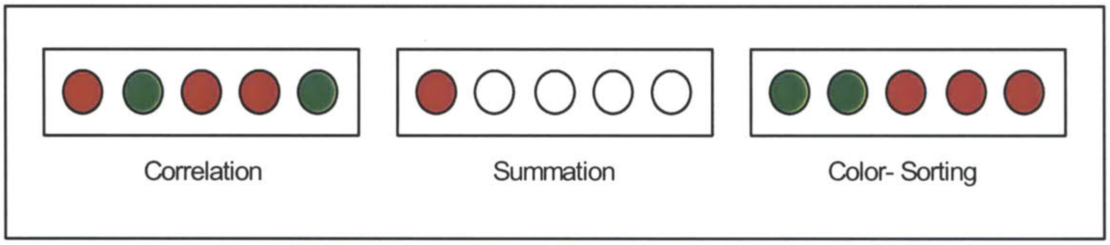

Figure 2-2 Three Choices for Thinking Tag Display Algorithm (Users agree on two questions, disagree on other three)

There is some evidence that the Thinking Tag displays successfully functioned as a kind of "talk". For example, at one event, we witnessed two coworkers—one more senior than the other—interact. Their tags flashed four red lights, showing that they disagreed on four questions. Instead of interpreting this as negative, however, the more senior colleague said, "Perfect. We are complementary to each other." On another occasion, a participant saw that he was a better match with two people at the same company than they were with each other. The consensus was that this was a very good thing, since in actuality, this person served as a liaison between these two people's groups.

Both of the above examples show how people were able to interpret the Thinking Tag displays in terms of knowledge they already had about the participants, and how the interpretations were made richer by this knowledge. This result ran counter to some people's intuitions that the tags would function less well in a community of people who already knew each other. It is our experience that the opposite is true, since people's understandings of each other provide a powerful resource for making sense of tag content.

As an example of how people used their commonsense knowledge to interpret their tag displays, consider that different people ascribed different valences to LED color. Most people seemed to subscribe to the notion that "Birds of a feather flock together" and wanted to find others with whom they agreed. They found green lights desirable. Others believed that "opposites attract" and went in search of five red lights. The fact that people were able to interpret these displays in the context of their personal beliefs further suggests the displays were treated like other utterances in the conversation.

2.2.2.2 Bootstrapping Conversation

When two people approach each other, before they say anything, their Thinking Tags begin a quick conversation of their own: one tag tells the other about its answers to the five opinion questions; the other responds with its own answers. Each tag then compares the opinions of the other tag to its own, and then finally displays the result of this comparison to the tag's viewer. This conversation protocol is very similar to the process Goodwin outlines, where "within conversation, subsequent utterances display participants' analysis of prior talk: 'it obliges its participants to display to each other, in a turn's talk, their understanding of other turns' talk.'" (Goodwin 1990) As a means of launching a conversation between users, the Thinking Tags displayed their understanding of their conversation that had already occurred. This allowed the Tag users to pick up the conversation from there.

2.2.2.3 Viewer Customization

Unlike traditional nametags, whose content varies only with their wearer, the content of the Thinking Tag also varied with the viewer. The interpersonal nature of the Tag's content was also similar to that of talk. As Goodwin states, "Talk, rather than being performed by an abstract, isolated speaker, emerges within particular speaker/hearer relationships… Speakers design their talk taking into account their particular recipient of the moment…"

2.2.3 IR Communication and Conversational Distance

In order to support casual conversation, the tags had to be able to communicate with their conversation partners at a distance that was socially comfortable for the Tag wearers. The challenge was to ensure that both the wearer of the Thinking Tag and the Tag itself were in agreement about who their conversation mate was.

We used infrared (IR) communication to determine a participant's conversation mate, since the line-of-sight nature of IR transmission makes it possible to tell whom someone is facing. We discovered that tuning the IR output power involved a trade-off. If the IR output power was set too low, people had to get uncomfortably close in order to get their Tags to interact, hindering their ability to converse naturally and comfortably. However, if the output power was set too high, there was too much opportunity for "cross talk" between Tags at a distance. This made it impossible to ensure the integrity of the link between the two people interacting and the affinity display on their tags.

Of course, appropriate conversational distance is notoriously gender and culture specific (Sussman and Rosenfeld 1982), making things more complex. We attempted to tune the IR to find a balance between these poles, settling on a power value that limited Tag interaction to about three feet. As it turned out, many people seemed to enjoy the deliberate gesture of pushing their Thinking Tags toward each other when they met – like a new kind of handshake – so keeping the IR low to transmit only across short distances worked well.

2.3 Designing for Resonance

Over several trials of the Thinking Tags, we discovered that one of the key factors in the success of the activity was the quality of the questions. If people did not find the questions meaningful, they did not feel their Tags represented them, and consequently they had little interest in how their opinions compared to others. We also discovered that authoring good questions was very challenging. We started to formulate a few heuristics for what made good questions. For example, we realized that good questions were tailored to the people at a particular event. Good questions were evocative, but not invasive. Good questions pointed to something larger; they weren't just about simple matters of fact. Finally, good questions divided the community among its answers. A few of our more successful questions, based on informal user feedback, were:

> What institution will be the last to give up centralized control? (asked at a conference on complexity theory)
>
> A. The Military
> B. The Classroom
> C. Your Family
>
> What was the hardest thing for you to learn? (asked at a workshop on learning)
>
> A. How to program a computer
> B. How to ride a bicycle
> C. How to share

By contrast, we found that questions that were explicit instead of evocative were not popular. For example, we also asked the following at the learning workshop:

> What type of learner are you? (also asked at workshop on learning)
>
> A. Visual
> B. Auditory
> C. Experiential

In retrospect, we now know that successful questions, like successful folklore, resonated with many members of the audience. That is to say, they were predicated on experiences, assumptions and values that people identified with, and their circulation allowed people to experience a shared understanding of these experiences. For example, for the question on learning, most people – especially those who have an interest in learning – have strong beliefs informed by personal experience about what is hard to learn. We picked the three answers to connect with a variety of beliefs and experiences the participants might have. People who felt that learning to ride a bike was the hardest may have associated this type of learning with the challenge of acquiring a physical skill, or with an emotional interaction with whoever taught them. Others who answered that learning to program a computer was hardest may have associated that with the frustrations of trying to think linearly. The idea was to create a question and answers that individual participants would feel were speaking uniquely to them, to their experiences and beliefs. Then, when people found others who shared their answers, and – by implication – their beliefs and experiences, they felt they shared something important with them.

Clearly, the last question above, about learning styles, failed to resonate with most people. This question lacked the connections to people's beliefs, experiences, desires, etc. Instead of allowing people to read themselves into it, it forced them to identify themselves with a narrow category that had no particular associations for them. Therefore, many participants resisted answering it.

Even though we were successful in creating many resonant questions for the Thinking Tag activity, we came to believe that the participants of a social gathering classifying themselves along five three-valued dimensions that were determined in advance by the event organizers was very limiting. We wanted to step out of the difficult role of designing resonant texts that would help people discover and experience commonality at a gathering. Furthermore, we wanted to open up the space to include many more texts that would have a better chance of finding more "pockets of resonance" within a large gathering. These limitations propelled us toward the Meme Tags discussed in the next chapter.

2.4 Events as Community Laboratories

Before the first Thinking Tag trial, we were very focused on making a new kind of computationally augmented nametag for a conference-type event. After the success of the event, however, it seemed clear that this type of technology might have relevance beyond conferences in more real-world, longer-term communities. That being said, we also came to believe that conference-type gatherings functioned well as a kind of laboratory for exploring technology that supports face-to-face community building. This "event as laboratory" model was important, because one cannot test this type of technology in a classic experimental setting, with a small number of subjects and a few hours of time. Events became the perfect middle ground between the sterility of the traditional laboratory and the intractability of the real world. The following paragraphs discuss the key reasons why events are good community laboratories.

2.4.1 User Community

At the Media Lab's tenth anniversary, we handed out Thinking Tags as sponsors entered the Lab. Newcomers could see that everyone else was already wearing the tags, and they had an immediate base of users to try their tags with. Also, sponsors knew they were at a special event that would only last for two days, making them more willing to try something that was unfamiliar.

Folk Computing technology requires a user community of sufficient size to be meaningful. It also requires a high "saturation" rate: users must encounter other users frequently enough to make the technology both useful and socially acceptable. Finally, it requires a community willing to try something new. These conditions are easier to establish at an event.

2.4.2 Maintenance

It is possible to sustain a new type of augmented social interaction over a two-day event in a single building that could not be sustained over a larger amount of time or space. The small "space-time footprint" makes it is possible for a reasonable number of dedicated support staff to keep the technology running.

2.4.3 Leverage

All the events where we deployed Folk Computing devices were already going to happen with or without our participation. At a minimum, this meant we could leverage the work of others on event planning and execution while we focused on the technology. More importantly, it meant we could leverage the "semantic field" that the event created. Because folklore happens in the wild – in the gaps between community structure – it wouldn't have made any sense to bring people together to try out some Folk Computing devices as a foreground activity. The people who came to the events we worked with comprised an authentic community, or at least a community-in-the-making, with an already established sense of what they were doing there. The event organizers had already worked to establish a meaningful structure for the event to unfold within. These structures then served as a trellis for the vines of folklore and Folk Computing to grow around.

2.4.4 Technology Development

One can often simplify the technology development process by taking advantage of the event's time-space constraints. Originally, we were going to deploy the Thinking Tags at an event in a large auditorium. As a precautionary measure, we took two early prototypes of the tags over to the auditorium to see how they worked in that environment. To our surprise, we discovered that sunlight streaming in through the giant windows flooded the infrared communication between the tags, making them inoperable. Rather than redesigning the communications circuit, however, we simply decided to deploy them at the smaller Media Lab sponsors-only event at a different location with little exposure to sunlight. We also took advantage of the event's time-space constraints in our design of the power circuit: we used the smallest batteries we could find that would power a Thinking Tag for the duration of the event.

From our perspective, we were taking advantage of the event's particular space-time properties to make the technology work. From the user's perspective, these properties seemed "natural" in the context of the event, and they treated the technology like it "just worked." This was good, to the extent that it let us observe how people might experience the technology in a natural setting. However, it also created a "Lost Horizons" problem when sponsors wanted to take the technology home with them. Something similar happened in Capra's classic film where a visitor to a Utopian land where no one grows old falls in love with a woman and steals her away to his homeland. Sadly, his hopes for a life with this woman are crushed when he witnesses her suddenly age and then die upon leaving Shangri-La. With Folk Computing, the technology is the woman, the Media Lab is Shangri-La, and there was disappointment when sponsors tried to bring the Thinking Tags home with them: their battery died almost immediately. More importantly, the tags seemed much less interesting outside the environment where there were many other people who could interact with their own tags.

3 Creating a "Folk Culture" Culture: The Meme Tags

The Thinking Tags proved that technology could play a positive role in supporting informal, face-to-face communication. Furthermore, they demonstrated how circulating, resonant texts could help people establish a sense of common ground. Experiments with the Thinking Tags revealed some limitations however.

One of the major findings of the Thinking Tag experiment was that the quality of the questions played a central role in the success of the activity. At every Thinking Tag event, however, there were people for whom most or all of the questions lacked meaning. Although we were pleased with the resonance exhibited by many of the Thinking Tag questions we authored, it was clearly a problem that the participants of a social gathering were only able to classify themselves along five three-valued dimensions determined in advance by the event organizers. Instead, we wanted to allow participants in the community to create an arbitrarily large and unfolding space of resonant texts within which they could locate themselves and others. This meant it would no longer be acceptable for participants to select which texts represented them all at once at the beginning of the event.

Another limitation of the Thinking Tag was that its five LED display that could only reflect patterns of resonance between two people. Although we were interested in the role the Tags could play in terms of integrating the community, many participants focused exclusively on the pair-wise "matchmaking" potential of the Tags. The most frequently asked questions about the Thinking Tag technology were variants of "Have you tried using this in a singles bar?" This aspect of the Tags received a lot of press attention as well (McCrone 2000), and some companies ultimately commercialized similar devices aimed at single people.

In order to address some of the limitations of the Thinking Tags, we needed to draw on a larger set of human communication constructs than those relating to the basics of face-to-face conversation. The Thinking Tags were designed in terms of some of the rules of informal talk. The Meme Tags utilized a well-established communication system that is "built on top of" talk, a system that affords a decentralized, democratic approach to the construction and circulation of resonant texts: namely, folklore. We wanted to design a technology that would substantially reduce the cost of participating in the folkloric process, while increasing the ability to visualize the patterns of resonance created by it. In this way, we aimed to create a "folk culture" culture: a social context that would encourage the creation of a folk culture.

This chapter starts out with a brief overview of the Meme Tag technology and activity. It then provides a rationale for Meme Tag design in terms of the two major Folk Computing objectives it addressed: reducing the barriers to participating in the folklore process, and visualizing the circulation of folk texts.

3.1 Technology and Activity Overview

We conducted the Meme Tag trial at a set of fall sponsor events at the Media Lab. Over the course of four days, we carried on an experiment with four hundred students, faculty and sponsors from the News in the Future, Digital Life, and Things That Think consortia. When participants at the event picked up their Meme Tags (see Figure 3-1), the Tags had their name and affiliation printed on the front (as well as programmed into the Tag), and they were already programmed with one or two memes. Each meme was a maximum of 64 characters long, taking up two 2 x 16 character "pages" on the Meme Tag screen, shown in quick succession. Memes used to "seed' the Tag were drawn from a collection collected from the community in advance of the event.

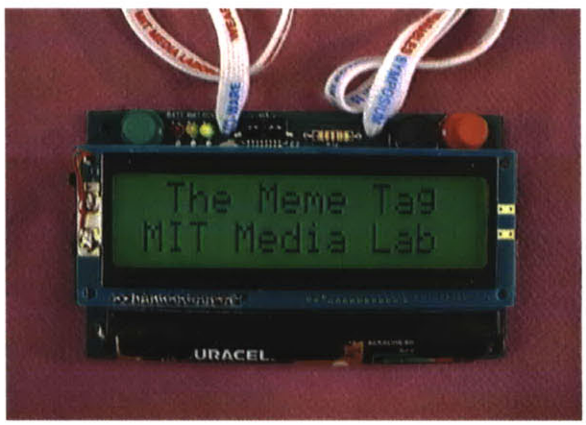

Figure 3-1 The Meme Tag

Once participants received their Meme Tags, they were free to roam about the Media Lab and exchange memes with fellow participants. For example, when Bob and Nancy meet and their Meme Tags activate, Bob's tag presents a new meme to Nancy, while Nancy's tag simultaneously presents a fresh meme to Bob. For example, Nancy's tag might say

```
Fresh meme for Bob:
  Computing should be about
  insight, not numbers
```

while Bob's tag displays

```
Fresh meme for Nancy:
  Make money fast --
  pass this meme to your friends!
```

If Bob likes the meme shown on Nancy's tag, he can press the green button on his tag, causing the meme to be replicated onto his. Similarly, if Nancy wants the message Bob's tag has shown her, she can capture it onto her tag. After their exchange, the Meme Tags become silent and do not distract from their subsequent conversation.

In addition to subscribing to memes from other people, participants were able to author their own memes at a kiosk and add them to their tags (see Figure 3-2). Around the event, large-screen displays presented visualizations of how the memes spread throughout the community. These displays formed a "Community Mirror," where participants saw in real-time which ideas were most popular and which ones were dying out.

Figure 3-2 The Meme Authoring Kiosks

3.2 Minimizing the Barriers to Participating in the Folklore Process

The Meme Tags were designed to decrease the personal and social costs associated with participating in the folkloric process. The hypothesis was that if this cost could be reduced, then a rich folklore could develop in environments that were not usually hospitable – for example, short-duration academic conferences. This computationally augmented folklore, along with the tracking visualizations it makes possible, would then help participants build a sense of community. The following sections discuss our attempt to simplify three aspects of the folklore process: authoring, performing, and recording. We then conclude with a discussion of whether we succeeding in lowering the cost of participation to such an extent that people did not have enough stake in the activity.

3.2.1 Reducing the Need for Immediate Authorship

Authoring memes was the most challenging part of participating in the Meme Tag activity. We were unable to design a method for authoring memes that felt well aligned with the dynamics of a social gathering. Also, many people – especially adults – do not think of themselves as creative, and we knew they might at first resist trying to author a 64-character phrase that their fellow participants would find compelling. We reasoned that once people had experienced trading and tracking memes, their willingness to author them might increase. Therefore, we looked for ways that people could start participating in the activity without first having to author their own memes. The key was to provide other means for acquiring memes on one's Tag. There were three ways of doing this:

Pre-installation: In the first Meme Tag trial, we initialized people's tags with three memes that were drawn from a pool of two hundred memes authored in advance by members of the community. We hoped this would allow users to begin exchanging memes immediately, and to experience the activity before having to take the time to author a meme. This choice had the unintended consequence of hurting people's identification with their Tags, however. We had assumed that soon after people received their tags, they would scroll through the memes on them, and delete the ones they did not like. Many people did not do this. Instead, they were unhappy when their Tags showed memes they supposedly endorsed – but had in fact not ever seen – to others. For the second trial, we reduced the number of pre-installed memes to one per tag, and introduced the "Poster Tags."

Poster Tags: The Poster Tags were designed to allow people to quickly acquire some memes they could trade with others, without sacrificing their identification with the contents of their Tags. In some ways, their function was similar to the buckets in the Thinking Tag activity. They consisted of a large piece of poster board with an imbedded Meme Tag, along with some printed instructions. The Poster Tags were positioned near the registration desk, so people could use them to add a few memes of their choice to their newly received Tags. The Poster Tags functioned like all the other Meme Tags: People simply approached the poster, and the imbedded Meme Tag would then offer them one of its memes. If they liked it, they would hit the green button on their Tag, and then get a copy of the meme. By allowing people to both practice using their tags and to control what memes went into them, the Poster Tags seemed to do a much better job of launching the Meme Tag activity than simply pre-installing memes on the Tags.

Trading: Trading was the final way to get a meme into one's tag without authoring it directly. One could participate quite fully in the Meme Tag activity by simply being a memetic intermediary, receiving memes from one set of people and passing them on to others.

3.2.2 Authorship

We considered a few ways that people could author memes without having to stand in front of a computer and cut themselves off from their social surroundings. One possibility was to have participants tell their memes to a "telegraph operator" behind a counter, who would then type the meme into a computer and download it into the user's tag. The challenge of creating a meme that was just 64 characters required more iteration than this type of interaction could provide, however. We also considered ways to make meme authoring more ubiquitous. For example, we contemplated modifying Dymo hand-held labelers so they could be used anywhere in the lab to program a new meme into one's tag (see Figure 3-3). This plan was discarded when we realized how frustrating it would be to author a meme this way, after the novelty had worn off. We ultimately chose a straightforward design for allowing people to author their own memes. We created kiosks with PCs where people could type their memes into a template, and then download them to their badges.

Figure 3-3 Possible Meme Authoring Device: The Dymo Labeler

We reduced some of the risk of authoring a meme by making it, by default, anonymous. As they circulated from badge to badge, memes did not reveal any information about their authors. Therefore, it was impossible to tell whether the owner of a Tag displaying a particular meme was the meme's author, or merely a carrier. We hoped that by knowing that their memes would not be discernable as their own, people would be more comfortable adding their own memes to their tags, and introducing them to the population at large

The Meme Tags further reduced the risk of authoring a folk text by establishing a special "communication channel" for such texts. Because Meme Tags "framed" their content in such a way as to remove any ambiguity about its folkloric nature, authors (and other propagators) did not have to worry about properly contextualizing this lore, with the attendant risks of misinterpretation.

3.2.3 Choosing What Lore to Perform

One of the significant aspects of the folkloric process is choosing what piece of lore to perform for a given audience. Choosing the right lore – the right joke, for example – for the right audience can be a source of satisfaction. However, people meeting at a conference setting do not often know each other well, and they might not feel comfortable or competent to make decisions about the appropriate lore to share. In other words, their lack of shared understanding makes it too risky to try to improve their shared understanding. This is a variant of the common ground problem discussed previously.

In an effort to solve the common ground problem, Meme Tags take over responsibility for choosing an appropriate piece of lore for another participant. We considered several algorithms for determining what lore to perform:

Fresh Memes: When two people met at the Meme Tag event, each person's Tag attempted to show the other a "fresh" meme to which the performer subscribed and the viewer had not yet encountered. The emphasis on freshness was to ensure that participants encountered something new that might help them start a conversation, and that memes would have the chance to circulate widely and trace the boundaries of their resonant community. This type of display did not work as a μ-cue because it did not reveal any common ground between the two participants. Participants had the chance to discover common ground in the same way that someone looking at someone else's T-shirt discovers that they're both fans of the same band. Unlike a T-shirt, however, the Meme Tags were able to show different content to different viewers, depending on what the viewer had already seen.

User-Chosen Memes: While conducting some preliminary tests before the Meme Tag trial, some people became frustrated they did not have more control over what meme they displayed to someone else. A suggestion was made that we give users full control over which memes got displayed when. While users might do this with their friends, we believed this placed an unreasonable burden on two strangers who were trying to get to know each other. As a compromise, an "expert" feature was added that would allow a user to offer a specific meme: by turning the knob on the tag to a specific meme, that meme would be offered in the subsequent exchange.

Collaborative Filtering: We explored collaborative filtering (Shardanand and Maes 1995) as a means for choosing the meme one person's tag would offer another. We hoped that such an algorithm would help identify memes that would be especially resonant with the recipient, and would consequently establish a pattern of shared resonance between the recipient and the sender (who is assumed to resonate with the meme to which she has already subscribed). Collaborative filtering algorithms make recommendations to users about instances of a particular type of object – for example, books – they might like. First, the algorithms identify others who like many of the same books as the user. The algorithm then can identify and recommend objects that are unfamiliar to the user but enjoyed by others of similar taste

We ran a small and informal test using collaborative filtering with Meme Tag-like objects in a virtual community. There were several benefits of experimenting in an online community, including: the lower cost of building devices in software instead of hardware; the opportunity to experiment with an existing, ongoing community; the ability draw on personal data from a central location to produce visualizations; and the ability to work with a subset of the community without being disruptive to the whole.

We built a Meme Tag-like object "worn" by about twenty guests for a birthday party hosted in the MediaMOO text-based MUD (Bruckman and Resnick 1995). When one guest addressed another for the first time, his Tag would display a meme chosen for that person via a collaborative filtering algorithm we implemented.

The results of the "Virtual Meme Tag" trial of collaborative filtering were not encouraging. We heard from several participants that they did not experience the choice of memes they were offered as anything other than random. There were two implementation factors that might have explained this. The number of people in the test group may have been too small for statistically robust collaborative filtering. Also, as implemented in a MUD, the Virtual Meme Tags themselves were somewhat awkward, which may have obscured the significance of the choice of memetic content.

We discovered one inherent limitation of using collaborative filtering in the context of Meme Tags that might have accounted for its lack of impact. To preserve the interpersonal quality of the meme swapping activity, a Tag can only recommend a meme to which its owner subscribes. Otherwise, there is no opportunity for creating a sense of shared resonance. However, this puts a large constraint on the pool of possible recommendations. It is very possible that none of the memes subscribed to by one participant would be deemed suitable for another via the algorithm. We realized that if we were going to build collaborative filtering into the real Meme Tags, we would need to build into the display a measure of confidence in the recommendation. Such a display would hopefully build users' confidence in the algorithm by keeping them from evaluating it when it was not able to function properly.

In the final analysis, we decided that the potential benefits of collaborative filtering did not clearly outweigh the costs of implementing it for the first (and only) Meme Tag trial. Implementing collaborative filtering on the Meme Tags would have been complex. The traditional implementation of the algorithm requires all available knowledge about each participant's predilection toward each object. This data existed on the Meme Tag server, but not on the Meme Tags themselves. We would have had to design a decentralized version of the algorithm that would have worked with a smaller sample of the total data available on each tag. It was unclear whether such an algorithm would be able to provide high-quality recommendations. Based on our experience with the negative value of low-quality recommendations, we decided our time was better spent elsewhere.

Rare Shared Opinions: Although the activity required that participants be presented with memes they hadn't seen before, we also considered combining this with an approach to select memes that would be familiar to both parties. The different approaches represent different ways of building common ground. In the case of presenting unfamiliar memes, the viewer is invited to join a community in which the performer is already a member – namely, the community of people who share some underlying beliefs that cause them to resonate with that meme. For example, if the meme is "It's not that I didn't think of it, I just didn't do it", the viewer discovers that he is not the only person who has had to deal with the frustrating experience of hearing countless suggestions on how to improve a demo that were obvious, but not feasible. In the other case, when the shared nature of a meme is revealed between two people, they learn that they are both members of that meme's community, and therefore share some underlying beliefs between them.

In order to enhance the value of revealing a meme two people share in common, the Meme Tags should choose the meme they share that has the smallest number of other subscribers. This derives from the basic tenant of Shannon's Information Theory (Shannon 1948), which suggests that the significance of two people sharing an opinion will vary with the inverse of the probability of this outcome, and the probability will vary with the percentage of the larger population that shares this opinion. For example, two strangers who discover they both believe the earth is flat will feel they have more in common than two who discover they both believe the earth is round. Similarly, two people who discover they share a disdain for a movie or book that everyone else has raved about will experience a stronger sense of commonality than two people who share the prevailing opinion.

For the Meme Tag debut, we chose not to implement a "meme in common" display. We were concerned that we would not be able to design a usable interface to allow participants to either switch modes themselves, or to understand which mode the devices chose at a particular time. However, we were subsequently able to informally test the efficacy of revealing "rare shared opinions", and the results were promising. Once again, in order to experiment, we took advantage of the unique affordances of an online community.

Our experiments involved the Foresight Exchange (www.ideosphere.com): an established community of several hundred people who place bets on the outcomes of various member-proposed predictions (e.g. "George W. Bush Remains President in 2004" and "Apple Computer Dies by 2005"). Because this community requires its members to be explicit and systematic about their opinions (in the form of bets made on a standardized set of predictions), it is an ideal test-bed for the computation and deployment of a variety of μ-cues, including rare, shared opinions (for more details about this play-money futures market, see (Hanson 1990)).

Using our Foresight Exchange proxy server, a user who reads a message on a discussion list automatically finds out about the opinion that he and the author share that is shared by the fewest number of other community members (of course, they may not share any opinions). In our experience of using this, we found this type of revelation was compelling, and that the new shared understanding that it brought was useful for interpreting that person's message. We also got feedback from a few FX users, who also felt this type of augmentation was revealing and useful.

3.2.4 Performance

In an effort to make it comfortable for people to share a piece of lore with someone they may have just met, the Meme Tags took responsibility for the performance of the lore. In the case of the Meme Tags, performance is defined in a fairly limited way: the display of a 64-character piece of text on an LCD display for an audience of one person. Nevertheless, there were several important design decisions relating to meme performance, including the decision about when the performance should take place, and whether there should be an introduction:

Timing of Performance: There were several issues broadly related to the timing of Meme Tag activation, such as: the right place in the conversation to activate the Tags; the appropriate duration for the Tag display to remain on; whether both Tags should be activated at the same time; and whether the Tags should reactivate when talking to the same person. Of course, good timing is the legendary hallmark of expert story- and joke- tellers, so appropriate choices in this area were important.

Using the metaphor of the third-party introduction, the Meme Tags were designed to do their work at the beginning of the interaction, and then disappear into the background. The Thinking Tags made this "disappearance" straightforward: participants could read those tags' five-LED display quickly and then move on. With their two screens of text, 32 characters each, the Meme Tag displays were more complicated. We had to experiment with the appropriate display time per screen, as well the number of times to cycle through the meme, in order to ensure the meme was displayed long enough to read but not so long as to cause a distraction. After some experimentation, we set the Tags to display the selected meme three times and to pause on each screen for two seconds.

The Meme Tag displays also needed to be timed in relation to each other. Ultimately, we designed them so two Tags in a conversation simultaneously displayed their memes for their respective viewers, in a manner somewhat analogous to two people starting a conversation by both telling different jokes at the same time. Obviously, this was a problematic choice. This choice also exacerbated the common ground problem, since – unlike with the Thinking Tags – participants could not tell what was on their own Tags by looking at the contents of their conversation partners' Tag. Instead, they were in the awkward position of having to rely on their partners' telling them what meme they were offered.

Unfortunately, there was no clearly superior alternative to the "simultaneous performance" approach. We considered having only one of the two Tags fire in an interaction, but then the person whose Tag activated would get no indication that he was part of an exchange. That problem could have been addressed by having the second Tag display the same Meme as the first, along with a note that says "FYI: This meme you currently subscribe to is being offered to other person". We believed the distinction between an "FYI" meme and a genuine meme offering would be unclear to many users, however.

The decision to display a different meme on each of the two tags simultaneously worked acceptably well. Participants, in general, were able to understand what was happening. In addition, it provided greater opportunity for memetic circulation, since in each interaction, two memes could replicate. However, additional experiments need to be undertaken to explore intelligible ways of allowing a single meme offer per interaction.

The final timing-related issue in the design of Meme Tag performance was deciding whether two Meme Tags should ever reactivate in the course of a sustained conversation between two people. Our initial design did not allow for this possibility, which we saw as potentially disruptive. The Tags were programmed to remember whom they last interacted with, and to not show another meme until they encountered someone different. However, in early experiments, several people said that sometimes they wanted to be able to exchange more than one meme in a conversation. Therefore, we added a feature that allowed participants to hit a button on their Tag to initiate an interaction with the next Tag seen.

Introducing the Performance: Research in the social roles of humor has highlighted the importance of introducing jokes via "prefacing devices" – pre-joke utterances that position the joke to a particular audience and signal the function it is meant to play (Cashion 1986). Since users did not usually choose the memes that would be displayed for others, they were not in a position to introduce them. Therefore, the Tags themselves needed to play this role by inserting some content before the meme.

The original design for the Tags did not include any meme prefacing device. The highest priority was to keep the meme performance short, which prohibited the display of an extra screen of non-essential text. In early experiments, however, we observed that it was difficult for viewers to be sure that a particular Tags' display contents were meant for them, and that they did not experience the content as being chosen for them. Since this greatly diminished the impact of the memes, we decided an effective introduction screen would be worth its cost in performance time.

The Tags took advantage of their very limited knowledge to formulate an introduction for a meme that would both personalize and position it. The result was a single screen display that popped up before the chosen meme, and read something like "Fresh meme for Mike…" Although the two-screen meme was displayed three times in rapid succession, as a means of saving time, the introduction screen was displayed only once. The word "Fresh" in the introduction was only displayed if the Tag was able to find a meme to which its wearer subscribed and its viewer had not yet encountered. The use of this word was meant to highlight the unfamiliarity of the meme, and more importantly, the fact that the performing Tag chose it because the viewer was not familiar with it.

The display of the viewer's name in the introduction turned out to be an extremely powerful demonstration that what viewers were seeing had been created just for them Of course, it is highly ironic that finding one's name on someone else's name tag is compelling. Conventional wisdom would say that this is the last thing anyone needs to see on someone else's name tag, since people know our their names. However, many people told us they found the Meme Tag prefacing device extremely engaging, almost to the point of distraction. One person said that if they walked past someone and their tags inadvertently started to interact, it was hard to resist stopping and talking when he saw his name lit up on the other person's tag. Perhaps that was why one of the speakers at an auditorium presentation during the trial concluded his speech and purposefully shielded his tag with his hand as he walked up the aisle toward the exit. He wanted to avoid the potentially awkward situation of his tag "striking up" an unwanted interaction with a member of the audience.

The Meme Tag personal salutation seemed to have the power to create what Goffman called a "focused interaction" between two people, which involves "individuals who extend one another a special type of mutual activity that can exclude others who are present in the situation" (Goffman 1963). In fact, we designed the Meme Tag software to ensure that one Meme Tag would seek out a single other one, and exclude other tags in the vicinity. We did not want one person's tag starting an interaction with several others' tags at once, leading to multiple meme offers and subsequent confusion about which meme might be accepted by pressing the green button.

Goffman describes an elaborate human protocol for negotiating focused interactions that includes such rituals as third party introductions of two people who have something in common. In some ways, the Meme Tags enacted this protocol by choosing two people in a group and lighting up their tags with memetic content drawn from one and personalized for the other.

Interpreting Tags as Performance: Despite our efforts to design the Tags so that they were effective at performing their memetic content, some people had difficulty with the concept of a wearable display designed primarily for the consumption of others. In fact, several Media Lab members demonstrated difficulty with this concept: in a brainstorming session on possible Meme Tag uses, they repeatedly suggested applications, such as using them as pagers, for which they were not well suited. Of course, Meme Tags are unusual in the world of technology, where pagers, PDAs, and cell phones that people carry are designed to be looked at primarily by the user.

Over time, however, as participants looked at the tags of the others, they seemed to get more comfortable with the idea that their tag was meant for others to view. In fact, this is a communicative model with which we are all familiar: namely, fashion. Fashion is about wearing things that communicate something about us to others (Davis 1992). Goffman talks specifically about the role of fashion in shaping one's "performance" in social situations (Goffman 1959). In some ways, the Meme Tag is another type of wearable display like a necktie, a piece of jewelry, or a T-shirt with text on it. Of course, the Meme Tag has the special ability to change its appearance depending on who is viewing it.

3.2.5 Choosing What to Lore to Adopt

The Meme Tags eased the need to author memes and took responsibility for their selection and performance. One aspect of the folkloric process that the Tags did not attempt to simplify was the selection of memes viewers chose to adopt as their own. In fact, we purposefully made this a little more complicated than it needed to be. For memes to trace out meaningful patterns of resonance as they circulated within the community, people needed to care about them before they subscribed. Therefore, we designed the Tags to ensure there was an opportunity cost to subscribing to a meme. We limited the carrying capacity of each Tag to seven memes. When users filled up their Tags, they had to delete a meme before they could accept any new ones.

The limited capacity of the Tags caused some frustration with users. It was particularly awkward when, at the beginning of a conversation, someone was presented with a meme she liked, hit her green button, and then received some annoying beeps and a display that told her conversation mate that her Meme Tag was full. If she decided to go ahead and choose a meme to delete, so that she could get the new meme, this further disrupted the conversation that the Tags were trying to support. Although it would have been easy to eliminate this disruption – for example, by automatically deleting the oldest meme in the Tag to make room for the newest – it was important to have people think carefully about what memes they wanted. Unlike the Thinking Tags, the Meme Tags were not solely trying to support pair-wise interactions. Therefore, we decided having a small percentage of disrupted conversations was a reasonable price to pay for satisfying one of the Tags' other goals: producing data for visualizations that reveal community-wide patterns of resonance.

3.2.6 Remembering Lore

Judging by the number of web sites and programs dedicated to teaching people how to remember jokes and stories (e.g., http://www.sharpsoftware.co.uk/total/ or http://www.dsv.nl/~tom/Menu_J.htm), the ability to remember and recall appropriate lore in a conversation is problematic for many people. The Meme Tags, with their perfect computational memory, remove this barrier entirely. Anyone wearing a Tag will be able to immediately and flawlessly recall a meme they had admired previously. Of course, this perfect recall eliminates the introduction of any variation that is part of the essence of folklore. This was a limitation we took up later with the i-balls.

3.2.7 Sanctioning the Folklore Process

Beyond reducing the cost of participating in the various aspects of the folklore process, the ubiquitous presence of the Meme Tags made it feel less risky to create and exchange folk texts. As a parallel, one can imagine that someone standing on a field full of people with baseball gloves would not feel self-conscious about pulling a baseball out and throwing it to someone. The very existence of the Meme Tags sanctioned participation in the folklore process. This was furthered by the existence of the Community Mirrors, which made this process even more visible to the community, and rewarded active players with a certain amount of fame (see below).

3.2.8 Results of Lowering the Barriers to Folklore Participation

Usage statistics suggest the Meme Tags were successful in lowing the barriers to participation in the folklore process. The data shows that 144 participants – 53% of the population who used their tags[8] – authored at least one meme during the two-day event. A total of 302 memes were created during the event, an average of about two per author. Although we have no control group to compare this to, it is hard to imagine a gathering without this type of technology producing as much lore. The data shows there was also a fair amount of folklore circulation: each participant accepted, on average, at least 4 memes from the tags of others.[9]

The data also suggest that the folklore medium created by the Meme Tags was usable and dynamic enough to respond to the unfolding needs of the participants to coalesce groups around particular beliefs. For example, after some conversations during the trial event, a Media Lab sponsor wanted to test the level of interest in a Europe-based sponsor gathering. Therefore, he introduced the meme "I'm interested in going to Euro TTT in '98" as a way to probe other participants at the gathering. It found five subscribers. The Meme Tag activity itself inspired several new memes, including observations like "Itispossibletoincludemanymorethoughtsinamemeifyouleaveoutspaces," and criticisms like "When did a meme become an aphorism?"

Unfortunately, other data led us to believe that we may have reduced the barriers to folklore participation so much that people no longer had a stake in the activity. At the same event where

---

[8]"Used their tags" means they hit their green button at least once to accept a meme during the event

[9] The actual number of accepted memes per participant was considerably higher than this. Unfortunately, an error in the Meme Tag code resulted in the loss of about half of the circulation data.

we did the Meme Tag trial, two other researchers created an "interactive poetic garden": a beautiful installation blending an actual rock garden and flowing water with projected text that appeared to flow along with the water current (White and Small 1998). They were looking for a way to relate the text that circulated in the garden to the viewers, so we collaborated on a mechanism to pull the text of memes off viewers' Tags and insert them into the stream. Although this worked well from a technological standpoint, they discovered that the Meme Tag/Poetic Garden connection was lost on many viewers who seemed unaware of what memes were on their Tags.

There are two reasons why people may have felt little connection to the content on their Tags. First, the resonant potential of 64 character strings of text may be inherently limited. This feeling was reflected in several meme contributions, including "`I weary of hunting aphorism`" and "`There isn't enough room to say anything of substance.`" It is also possible that we made it so easy for people to participate in the folklore process that they were hardly participating in it at all. By design, it was the Meme Tags – which selected which texts to perform, performed them, and also recorded the performances of others – which participated more fully in the process, leaving users with the modest job of pushing the green button when they wanted to accept a particular meme.[10] To paraphrase Richard Dawkins, Meme Tag users may have become a meme's way of making another meme (Dawkins 1989). This view was reflected in some of the memes that were contributed, such as "`Please let me propagate- Press that button`" and "`Your Meme Tag is broken. Press the green button now.`"

While the concept that the Tags were using the participants, rather than the other way around, may be compelling for sociobiologists, it is not a good model for community building. Indeed, we heard from some users who felt their tags had too much power over them. They did not want their tags to choose what meme would be displayed to someone else. The optional "manual override" we built in to let people chose what meme their Tag would offer did not satisfy these people's concern that their tag would offer an embarrassing or inappropriate meme to someone they encountered.

With the i-balls, the next generation Folk Computing technology, we worked to address both potential reasons why people felt disconnected from the folklore process. We enabled the

---

[10] Of course, participants also authored memes, but this was a relatively small part of the overall Meme Tag experience

creation and exchange of richer objects than short text messages. We also gave people more control over the process.

3.3 Community Mirrors

In the literature, folklore is often spoken of as a kind of mirror or reflection in which groups of people can examine themselves. For ethnic groups, Oring talks about folklore's "self-reflective role that allows the dynamics of the ethnic group and aspects of ethnic identity to be reflected back to its members" (Oring 1986). Dundes claims folklore serves "as a kind of autobiographical ethnography, a mirror made by the people themselves, which reflects a group's identity…" (Dundes 1989). Finally, Boas asserts, "folklore constitutes a mirror of culture" in (McDowell 1979). What exactly is reflected in this mirror, however? As defined by most folklore scholars, the mirror of folklore reflects the values, beliefs and experiences underlying the lore back to group that shares it.

Solid data about the distribution of lore within a community is missing from the reflection seen in the folklore mirror, however. While interpersonal interaction can establish a piece of lore as shared within one's local social group, people do not have good intuitions or data about where a piece of lore has circulated (see Section 1.2). In this way, the folklore mirror reflects more about the "why" of a larger folk group, and little about the "who".

Traditionally, it has been the difficult work of folklore researchers to combine an understanding of a set of folk texts, the underlying beliefs and experiences they reflect, and the meticulous effort required to uncover their temporal and spatial patterns of circulation. For example, Figure 3-4 depicts the distribution of a particular Estonian folk tradition regarding a type of ghost called a "home wanderer" (Viluoja 1996). Because of the amount of work required to gather and interpret this information, a view of folklore circulation like the one in Figure 3-4 is only available to a researcher; it is not something usually afforded to members of the folk group.

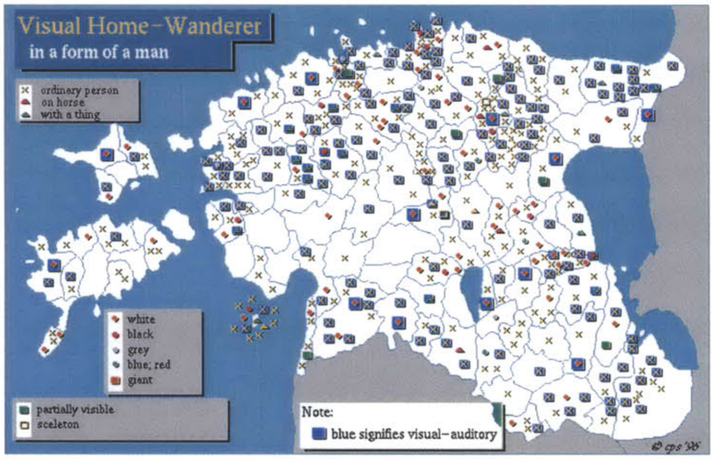

Figure 3-4 Folklore Researcher's Analysis of Geographic Distribution of a Particular Folk Text

Our goal was to produce in real time simplified versions of the kinds of visualizations that used to take folklorists months or years of effort and analysis, and to produce them for the immediate consumption of the groups they represented, not for an audience of fellow researchers. These visualizations were meant to give people a better understanding of the folk groups in which they were engaging, both in terms of those groups' constituent beliefs and members. This was made possible by the readily traceable nature of digital objects circulating among computational devices. The next section describes how we designed the technology to implement these visualizations.

3.3.1 Designing Technology to Track the Face-to-Face Circulation of Folklore

The technical challenge was to move transaction data from the Meme Tags to a central server where they could be transformed into visualizations. A transaction record consisted of a unique transaction id, the id of the Tag that offered the meme, the id of the Tag it was offered to, the id of the meme, and the code for what action was taken (i.e. was the meme accepted or rejected). One choice for moving the data was radio frequency (RF) communication. Since infrared (IR) communication only works well for line-of-sight communications, it was not well suited to transmitting data from a Tag to a server a crowded room. However, at the time of the Meme Tag trial, there were no readily available RF solutions that addressed our requirements for small size, low power, low range, low bandwidth, and low cost. [11] Furthermore, the Tags required IR and its line-of-sight quality to determine whom someone was facing (as a proxy for whom they

[11] Solutions such as BlueTooth are now starting to address this need (see www.bluetooth.com)

were conversing with). Implementing both IR and RF communication on the Meme Tags seemed prohibitively complex and expensive.

We conceived of a way to use IR for both person-to-person and person-to-server communication, by taking advantage of some of the properties of a conference gathering. We knew that at such a gathering, we could count on three things:

- Critical Interaction Density: People would engage in a relatively large number of interactions per hour.
- Key Gathering Points: There would be certain places where large numbers of people would be likely to pass by, such as the coffee and food tables.
- Bounded Geographical Scope: The gathering would be concentrated on the lower two floors of the Media Lab building.

With these constraints in mind, we designed a system where interaction data would travel from individuals to a server via an arbitrary number of intervening person-to-person hops. This type of communication system later came to be called a Mobile Ad-hoc Network, or MANET [A, 1999 #99], but at the time, we could find no other examples of this type of network to draw on. We drew some inspiration from a system designed to replicate information across a wired network in a viral fashion (Chesley 1991). Figure 3-5 shows how the whole system worked: when two people met each other, their Tags first performed the requisite communication to select an appropriate meme to offer. After the users had made their decision whether to accept the meme being offered, their Tags created a record for the transaction, and added it to their transaction store. Then, for as long as the two Tags remained in IR contact, they copied transaction records between their respective stores. In this way, Tags carried information about more transactions than they were actually involved in, and particular transactions were recorded on multiple Tags.

When users went to kiosks to author new memes, all the transactions on their Tags were downloaded via IR and sent to a server over a wired network. In case participants did not visit the kiosks regularly, we placed several IR data receivers in areas participants tended visit regularly: specifically, near the food. As a final precaution, we wrote software for a special type of Tag that would forgo the normal meme exchange and immediately start replicating another Tag's transaction store. If the other systems weren't working, we planned to don these Tags ourselves, walk around the event invisibly gathering transaction data, and then go over to a kiosk to have it uploaded.

Since it was not possible to test our networking algorithm with hundreds of users, we built a simulation. We used this simulation to tweak the algorithm's key parameter, which determined how many times a transaction record would get copied between two Tags. If this parameter were set too low, a transaction record would not be spread across enough tags, and might not find its way to the server. If it were set too high, the system would get flooded with copies of a small subset of transactions.

At the Meme Tag user trial, the technology that moved data from the Tags to a server performed marginally "above threshold." Unfortunately, a subtle bug in the Tag software caused transactions to stop being recorded in the transaction store after a certain amount of activity. This resulted in the failure of about half the transaction data to make it to the server. This was not a fault of the viral algorithm. In fact, the algorithm seemed to work well on the data that got into the system, and the Community Mirrors were able to display some visualizations that people found interesting.

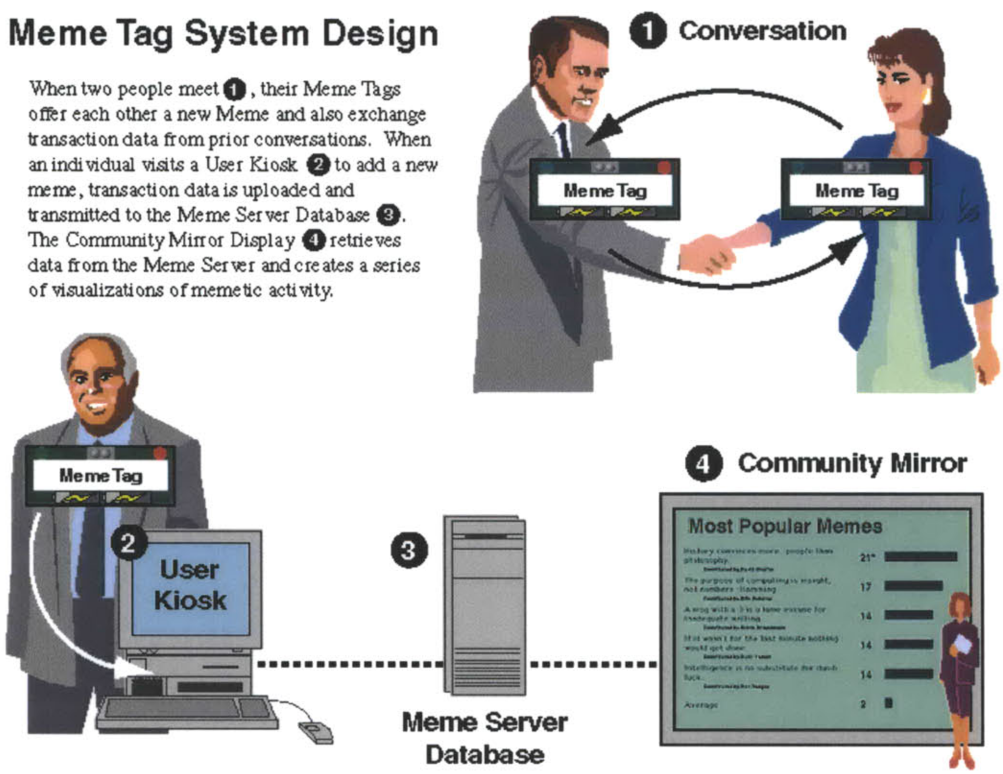

Figure 3-5 The Meme Tag System

Drawing by Fred Martin

3.3.2 Designing the Community Mirrors

The most basic design for a Community Mirror was a visualization that revealed to the group the memes that were the most widely shared at a given time (see Figure 3-6). The hypothesis was that these memes would reflect important shared values, beliefs, and experiences of the community. By establishing these beliefs as mutual knowledge (i.e. shared and known to be shared), they could form the basis of an interpretive community and a shared identity that would help participants communicate.

For example, one of the memes that showed up in the "Most Popular Memes" community mirror at some point was "`Your next computer is in your coffee cup.`" Professor Neil Gershenfeld, the head of the Physics and Media group, authored this meme as an in-joke reference to his Quantum Computing research project. In order to resonate with the prediction, one had to know that Gershenfeld was working on using the quantum mechanical effects of an amount of liquid small enough to fit in a coffee cup to do large amounts of computation, and one had to care. By reflecting the popularity of the meme back to the participants in the trial, the Community Mirrors established this knowledge and interest as shared, thereby making it available to use in the formulation of one's own communication and in the interpretation of the communications of others.

As μ-cues, the goal of the Community Mirrors was to transform distributed knowledge, where a group of people shares a belief, to mutual knowledge, where a group of people shares a belief and knows that the individuals share it. We used the principle of physical co-presence to establish the contents of the Community Mirrors as mutual knowledge, articulated by Krauss and Fussell: "If you and your friend were physically co-present at some event (and mutually know this), you could assume that the salient aspects of that event were also part of your common ground" (Krauss and Fussell 1990). Therefore, the challenge became to insure the Community Mirrors' salience. We did this by making the displays large (about ten feet diagonal), by centrally locating them in the main common spaces, and by making them call some attention to themselves via fast-changing, brightly colored visualizations.

Our technology to allow a gathering of people to reflect on their shared beliefs bears some similarity to real-time audience polling systems (1998)– the most famous being the one used on the television show "Who Wants to Be a Millionaire?" However, audience polling systems require more formal and centralized structure. The event organizers determine which questions will be presented to the audience, and the questions must be asked and results analyzed in the context of a formal presentation. With the Meme Tags, anyone could contribute a meme, and the collective activity of all the members of the group – played out in the exchange of memes in informal, face-to-face interactions – determined which memes got the attention of the group.

There is another key difference between the Meme Tags and an audience polling system that makes the Tags more appropriate for informal gatherings. Effective polling questions provoke a strong positive or negative response, such as "strongly agree" or "strongly disagree" (Jason 1997). However, the individual response to most memes (and to much of folklore) tends to vary from "strong resonance" to "failure to resonate." Folklore resonates with those who share its presuppositions but does not usually engender a strong negative response in those who do not (there are exceptions, such as ethnic jokes). This is what makes it an effective social probe in informal situations. The polarizing questions used with audience polling systems are more appropriate for that context, where the person who asks them has some authority, and the people who answer them have anonymity.

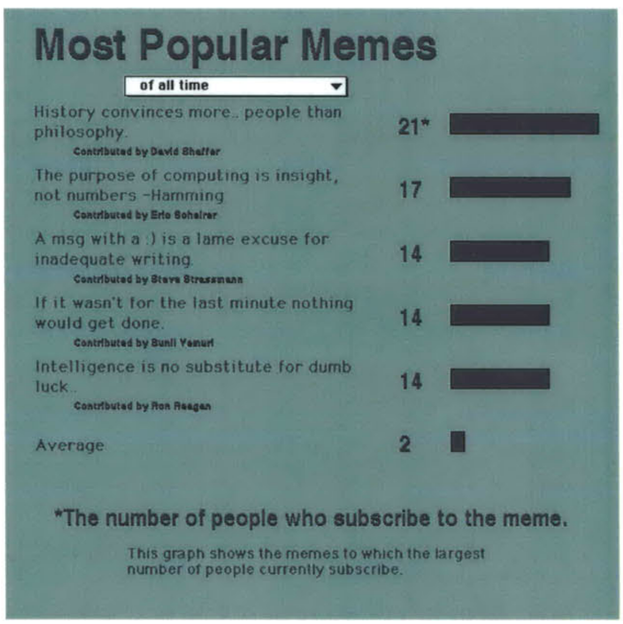

Figure 3-6 "Most Popular Memes" Community Mirror

There was a difficult trade-off involved in designing the Community Mirrors that displayed particular memes. We wanted the Tags themselves to create the distributed knowledge, by introducing people to new memes in the decentralized fashion highlighted above. Unfortunately, it was hard to design Community Mirrors that created mutual knowledge without creating some distributed knowledge as a side effect. That is what would have happened if people looked at the Most Popular Memes visualization who had not seen some of those memes before. The whole process would have been short circuited: the viewers would not have gotten to choose whether those meme resonated with them, that resonance would not have been factored in to the visualizations, and the viewers would not have had the option of adding the meme to their Tag.

In an effort to avoid affecting the distributed knowledge space, we originally designed the community mirrors to reveal only the first line of a two-line meme. We reasoned that this would be enough to recognize for people who had already encountered the meme, but not enough to give it away to those who had not. This turned out to be a poor design choice, however[12]. Many people found the displays confusing. Also, the visualizations were less effective "advertising" for drawing bystanders into the activity. Finally, it was not clear that because users saw a meme first on the Community Mirror, they would not still be interested in it when it appeared on someone's tag they were speaking with. If the goal was to ensure that memes had the chance to circulate informally before being revealed as popular, we could have also built in a time delay or a minimum threshold of popularity before a meme was displayed.

Even with its problems, the Most Popular Memes visualization was reasonably effective. Some of the memes on there seemed like they captured something basic about the Media Lab community in which they circulated, such as "`The purpose of computing is insight, not numbers`" and "`If it weren't for the last minute, nothing would get done.`" One faculty member reported a conversation between himself and a sponsor, where the sponsor kept looking over at the Community Mirror out of the corner of his eye. Suddenly, when the faculty member was in the middle of a sentence, the sponsor whirled around and took a picture of the display. He then explained that a meme he had created was on the most popular memes list, and that he wanted to document this.

The model of community we are trying to support is not one homogenous group that shares a set of beliefs, values and experiences, however. Nor is it a set of smaller, distinct and insular groups that communicate only within their borders. Instead, we are trying to support a community of multiple and overlapping "folk groups" at a variety of scales, with each group able to establish a sense of itself by reflecting on its own patterns of folklore circulation in the context of the larger community. Although our approach to this got more sophisticated with the i-ball technology (see Chapter 4), we first started exploring the concepts with the Meme Tag Community Mirrors.

---

[12] Community Mirror visualizations shown in this thesis are the modified design that shows both lines of a meme.

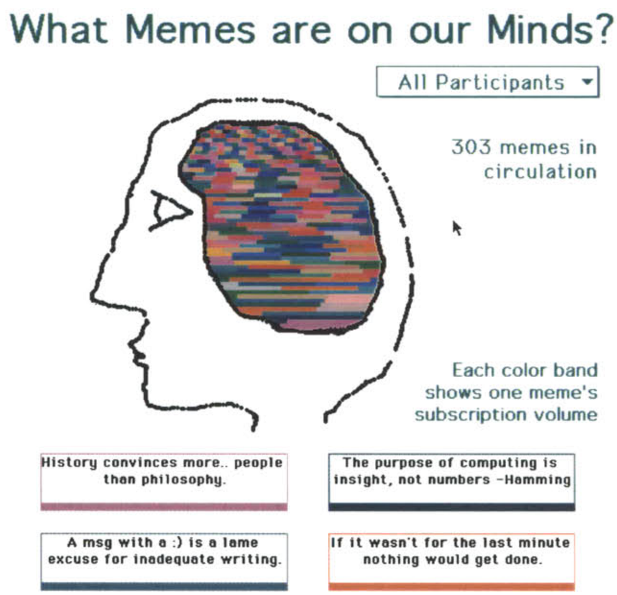

Figure 3-7 "Group Mindshare" Community Mirror

Toelken states that a key condition for a folk group is having "some factor in common that makes it possible, or rewarding or meaningful, for them to exchange informal materials in a culturally significant way" (Toelken 1979). To that end, we designed a visualization that showed the memes that were most popular within salient demographic categories – such as females or students – that were distinguishable by available data. Figure 3-7 shows a visualization that attempted to represent the collective "mindshare" of a particular viewer-selectable group. Each colored pixel in the brain corresponds to one of the seven available slots for a meme on one of the Tags of the members of that group. Each band of color corresponds to the number of people in that group who are carrying a particular meme (if the whole group was carrying the same seven memes, then the visualization would show a brain filled with seven large patches of color). The boxes at the bottom show the text of the four memes with the largest amount of mindshare within the group. Although the basic idea of exploring the popularity of a meme within a demographic group had merit, the particular Mindshare visualization had some problems. First of all, few memes gained significant mindshare during the Meme Tag trial. This was one of the findings that led us to design a construction environment for folk objects with greater expressive potential and a folk process that was more involving for the i-ball trial. Also, assembling all the memes in which a group had any interest made for a visualization that was hard to interpret. Finally, the Mindshare visualization did not allow for the comparison of meme distribution across different demographic groups.

In examining the data from the trial, we realized a more revealing approach might be to examine a particular meme's popularity across a demographic dimension (e.g., gender). For example, in the case of the meme "`This meme good for one free diner with Prof Tom Sayers`" (name changed), 9 out of 18 (50%) of the men who saw this meme subscribed to it, whereas only 1 out of 9 (11%) women did. This may have reflected widespread opinion that Prof. Sayers was sexist.

Another interesting example of a meme whose pattern of circulation revealed a distinction between two demographic groups was "`Nothing that is worth knowing can be taught.`" This meme resonated with Media Lab faculty and students, but not with sponsors: 35% of students and 56% of faculty who were exposed to it adopted it, but only 4% of sponsors did. It is possible to explain this resonance in terms of the find-it-out-for-yourself attitude that permeates the Media Lab, where teaching and course work are valued less than innovative research. It also makes sense that this attitude would not be shared by sponsors, who specifically come to the Lab to "be taught" about new technologies and their implications.

The above examples show how meme circulation can serve as a "marker of group difference," one of the characteristic social uses of folklore (Oring 1986). In our work on the i-ball project, which is the subject of the next chapter, we focused on building visualizations that could make these differences more salient than they would be during participation in the traditional folklore process.

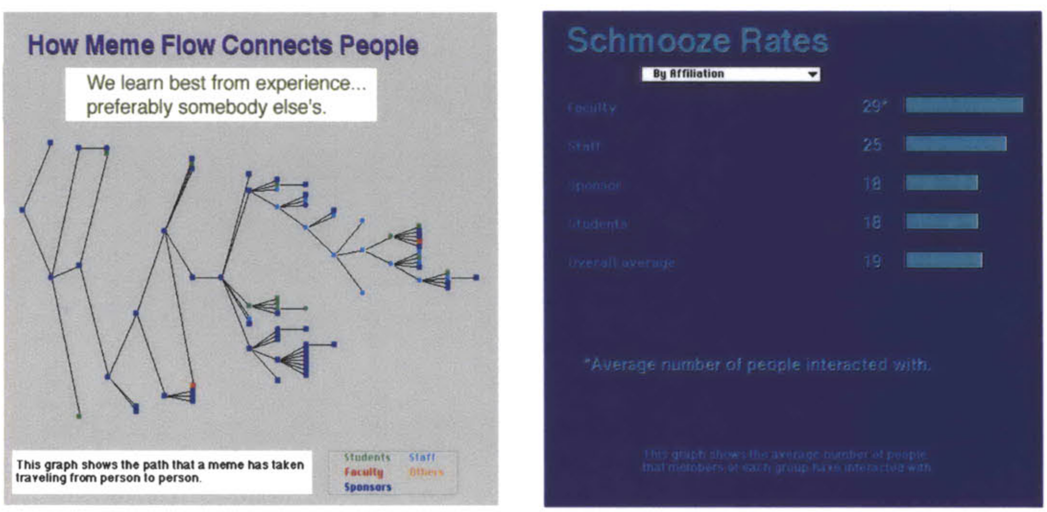

Figure 3-8 "Meme Flow" and "Schmooze Rates" Community Mirrors

The final type of visualization we explored for the Meme Tag project focused on the "social network" characteristics of the participant population. Memes did not just appear on people's tags, but rather got there through an unfolding series of interactions between pairs of individuals. These visualizations highlighted patterns in this interaction data. For example, the right-most visualization in Figure 3-8 ranked groups with various affiliations to the Media Lab (e.g. faculty, sponsor, staff, and students) by the average number of interactions they had participated in thus far. A more interesting (although perhaps less entertaining) visualization was a tree diagram that represented how a particular meme spread through a population. We colorized the nodes according to affiliation, which, upon close inspection, revealed a tendency for people to interact with others of the same affiliation. Because we found this type of insular behavior was interesting to participants, we worked to make it more perceptible in the next generation visualizations we built for the i-ball activity.

4 Designing an Environment for the Construction and Circulation of Resonant Objects: The I-Balls

The Meme Tag trial revealed the potential of allowing people to create, share, and track their own resonant texts as means for probing group similarities and differences within their community. Based on some of the results of the Meme Tag trial, however, there were two things we wanted to explore in our next Folk Computing technology. We wanted to allow users to create objects that were more expressive, but could still circulate through informal, face-to-face interaction. We also wanted to design a new set of visualizations capable of highlighting which groups resonate with which texts, and how particular texts move through a population. The result was the "i-ball" – short for "ball of information."

The following bullets summarize the difference between i-balls and the previous two Folk Computing initiatives:

- I-Ball objects had graphical and behavioral dimensions, as opposed to the purely textual quality of both the Thinking Tag questions and the Meme Tag memes.
- Existing i-ball texts could be easily "mutated" by participants to create new texts
- Users played a more active role in the exchange of texts. An exchange only happened when one person chose to offer a text and another chose to accept it.
- I-Balls were introduced into an existing community for an extended period of time – a few weeks instead of a few days

This chapter starts with an overview of the i-ball technology and activity, followed by a design rational for the i-ball construction environment. It then explores how we enabled the face-to-face exchange of i-balls and concludes with an exploration of the i-ball visualization tools.

4.1 Overview of Technology and Activity

The name "i-ball" is short for "ball of information." I-balls are simple software folk objects that have some toy- and game-like qualities. People can design their own i-balls and then share with other members of community. In our prototype, i-balls exist on key-chain-sized video game devices made by SEGA and sold as part of their DreamCast video game system (see Figure 1-3 in Chapter 1). We wrote our own software for this commercial device, and renamed it the "i-socket," to distinguish its capabilities from those of the original SEGA "Visual Memory Unit (VMU)"– designed to let kids store and recall their state in a particular DreamCast game.

Before the activity starts, participants input some standard information about themselves into their i-sockets. For example, at the school trial, students inputted their gender, their grade, their height, their favorite subject, the part of town they lived in, etc.

People then design their i-balls on a PC using a prototype graphical programming tool we developed. The most basic form of i-ball consists of a single "animation" programming block. This block allows kids to "decorate" their i-ball by composing an animation out of 128 different letters and icons in a simple "flip-book" style animation editor. I-balls created on a PC can then be downloaded to an i-socket via a small "docking station".

Like the play objects they are named after, i-balls can be passed between people. Participants can give a copy of one of their i-balls to someone else, or, using "jump" blocks and "rule" blocks in authoring environment, i-balls can be programmed to "bounce" from one person's i-socket to another's, based on user-defined rules – e.g., only jump if the person is of the opposite sex. I-ball passing can occur when two people connect their i-sockets using the ports that SEGA designed in to the hardware.

We pilot tested the i-ball technology in the fall of 1999 at a multi-generational conference on "learning through invention," where almost all of the 500 children and adults in attendance were given an i-ball device to wear around their necks. In the Spring of 2000, we did a more extensive trial over three weeks involving the entire third through eighth grades of a public K-8 school – 350 students, teachers, and staff members (including the principal, secretary, cafeteria monitor, etc.)

A few words about how this new technology was received and what kids did with it: At the first i-ball event, several people commented on the frenzy that the activity engendered. We had set up seven i-ball programming kiosks, and they were in constant use, usually with several people waiting behind each one. In a day and a half, more than 1600 i-balls were created – an average of more than three per attendee. The sight of participants passing i-balls from i-socket to i-socket was ubiquitous.

Over the course of our i-ball trials, children have created i-balls in a variety of "genres", including:

- "Hot potatoes" that must be passed according to certain rules in a certain amount of time

- "Quests", or scavenger hunts, that send a person in search of a variety of other people who meet particular descriptions – e.g. "find someone taller than you"
- "Randomizers" that were used to create "Magic 8 Ball"-type fortune teller toys
- "Hitchers" that are simple autonomous software "agents" that hitch-hike around the community, invisibly jumping from one user to another
- "Secret i-balls" that would show one animation to an "insider" group (e.g. students) and another when viewed by everyone else (e.g. teachers).
- "Multi-author i-balls" that children would add their own piece of animation to and then pass on

4.2 Supporting the Construction of Resonant Objects

4.2.1 Animations

The central element in the i-ball construction environment is the animation block, which allowed users to construct a multiple-frame animation out a fixed set of icons. In fact, I-balls consisting of a single animation block (*not* a single animation frame – a block is like a flip-book that can contain many frames) comprised 50% of the over four thousand i-balls created at the school, suggesting that this element in itself was capable of considerable expressive range. **Figure 4-1** shows the sequence of 27 frames that were a part of one such i-ball, called "Romance," created by an MIT undergraduate.

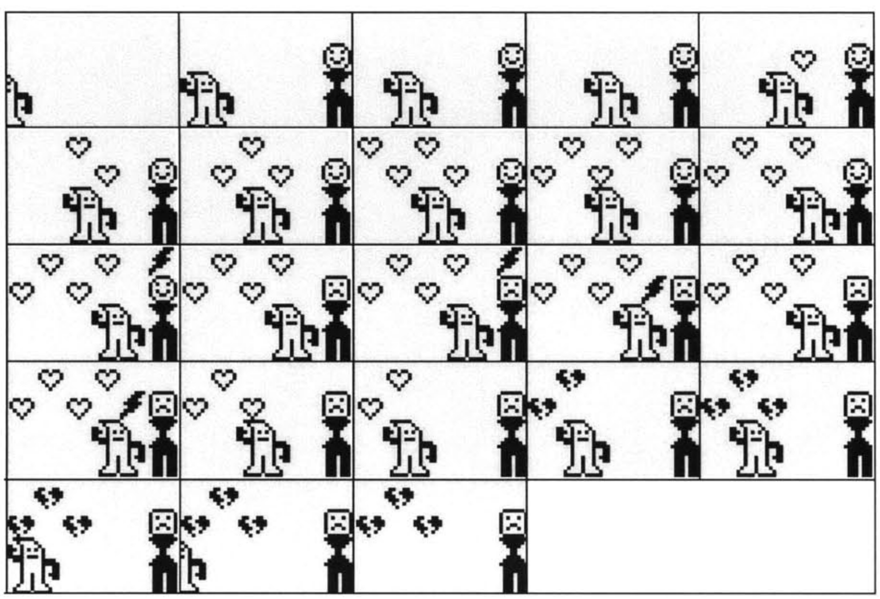

Figure 4-1 Animation Frames from Romance I-Ball (Read left to right, top to bottom)

The popularity of the Romance i-ball, among others, suggests that people more strongly resonated with i-balls than with memes. More than half of the participants at the school (131 of 251) got a copy of it. Compare this to the most popular meme in the Meme Tag trial, for which less than 50 people out of 316 participants had a copy.

Although the i-ball animation editor was relatively simple, users were able to create animations with enough subtext that their resonance varied with different subpopulations. This made i-balls effective probes for seeking out subpopulations with shared beliefs and experiences. For example, in the case of Romance, the data suggests that it resonated more strongly with females: 56% of all female participants got a copy of it, while only 48% of the male participants got it. One could attribute this to the tacit assertion of "girl power" represented by the female character, who forcefully rebuffs Gumby's amorous advances, pushing him off the screen into oblivion.

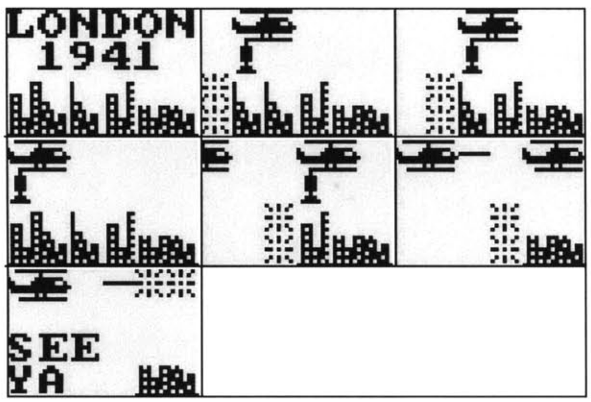

Figure 4-2 Animation Frames from WWII I-Ball

Figure 4-2 shows another example of an i-ball whose resonance skewed along gender lines, only this one was more popular with males. A male sixth grader authored "WWII". 43% of the male population got a copy of the WWII i-ball, whereas only 18% of the female population wanted it. This pattern of resonance may reveal something about the male population's bias toward violent content. It might also reflect men's satisfaction with visual effects. This was the first instance of the exploding cityscape motif that was used again and again. Women, on the other hand, may have been bored by the lack of narrative detail.

So that everyone could create animations to seek out likeminded subpopulations, we designed the animation editor in the middle ground between ease-of-use and expressiveness. To that end, we decided to use icons as the primitives for drawing animation frames, rather than letting people create drawings using more traditional drawing tools. Figure 4-3 shows the i-ball authoring environment with Romance loaded up. The set of icons that can be included in an animation are arranged in a rectangle in the lower right. In addition to the letters and numbers, we chose symbols and shapes that could be part of a wide variety of pictures. While the use of these icons as drawing primitives makes it easy for people to build representational images, there is no question it limited their expressive range. This range was limited still further by the particular icons we chose to include. For example, not including characters that were identifiably female, to complement LEGO Man and Gumby, was a glaring oversight. Many people also asked for the ability to draw at least a few of their own icons to add to the set. This would be useful functionality to add.

Another design choice that made creating animations relatively straightforward was the use of a flipbook-on-a-grid metaphor. Users drag icons from the palette onto a flipbook page laid out on 4 x 6 cell grid (in the lower left of Figure 4-3). When they have created the frame they want, they click on the "Add" button. A new page in the flipbook is created and the contents of the last page are copied onto it. Now the user can simply adjust the desired icons to create a sense of motion from frame to frame. With these simple tools, people were able to make a wide variety of entertaining and provocative animations.

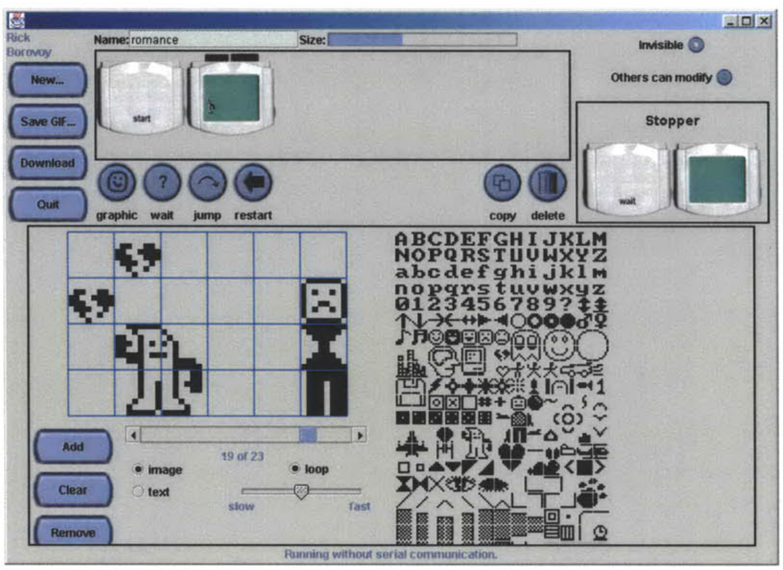

Figure 4-3 I-Ball Authoring Environment

4.2.2 Wait Rules

In addition to creating simple animations with the authoring tool, kids could create i-balls that exhibited more complex, interactive, game-like behavior. The key primitive that made this possible was the "wait rule block" – a block that paused the execution of the program until a particular condition was met. The following examples explore how we designed the wait rule block, and its role in creating objects that have more potential resonance than simple animations.

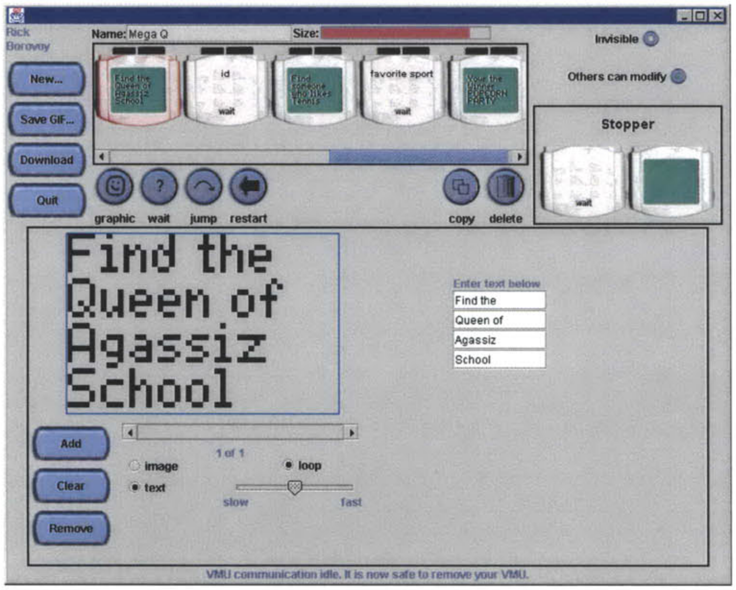

Figure 4-4 Authoring the MegaQ Quest

One of the most popular i-ball programming genres was the "Quest". Over the course of the school trial, 128 participants authored 287 different Quests. Quests resemble the classic "treasure hunt" children's play activity: get a clue telling where a particular object is located, then go find the object, along with another clue about the next object, etc. With a Quest i-ball, however, the treasure is another person. A clue directs the possessor of a Quest to find someone of a particular description. When the Quest holder finds that person, and mates her i-socket to his, a new animation plays, and the quester gets a new clue (or an animation that celebrates her finishing the Quest). For example, someone might author a quest that directs the player of the quest to find someone in the third grade who walks to school. When this is done, the player gets a new clue instructing her to find someone who likes broccoli.

The important role of the quest in literature and in children's games suggests that the concept itself resonates with people's experiences and desires. Authoring an i-ball Quest, in particular, gave participants a sense of power by allowing them to control the interactions of their peers. An i-ball Quest invited the people who played it to see themselves as on an important and mysterious mission. To explore in greater detail how i-ball programming allows for the construction of more resonant i-balls, such as Quests, the following paragraphs examine a particular i-ball in some detail.

Figure 4-4 shows part of the program for the "Mega Q" quest, an i-ball that was authored by the school's technology coordinator that became one of the three most popular i-balls of the trial. The i-ball program is the series of blocks toward the top of the editor. The left-most animation block is selected, and therefore highlighted in pink. Its contents are displayed below in the workspace area: a single text frame that says, "Find the Queen of the School". We built the text component of the animation block to allow people to create text-only frames by typing, not by dragging characters from the icon palette to the grid. When the i-ball is run, this piece of text will show up on the i-socket screen.

The next block to run in the "Mega Q" program would be the white wait rule block to the right of the animation block. If the user were to click on this block in the editor, he would see the contents of Figure 4-5. The workspace at the bottom of the editor shows the contents of this rule: "If all of these are true… your id is equal to 232" This wait rule will execute immediately after the preceding animation block is displayed. Because the rule says, "*your* id is equal to 232," the i-ball will wait until the quester connects her i-socket to another i-socket whose user id is 232 – ostensibly the "Queen of the School." (If the rule were written as "my id is 232", this would test the id of the i-socket on which the i-ball is currently running). When the user connects to the queen, the animation block to the right of the wait rule will execute and provide the next clue. If the quester connects to someone with an id other than 232, or does not connect to anybody, the wait rule will keep waiting. While it does, the previous animation block will continue to show on the screen.

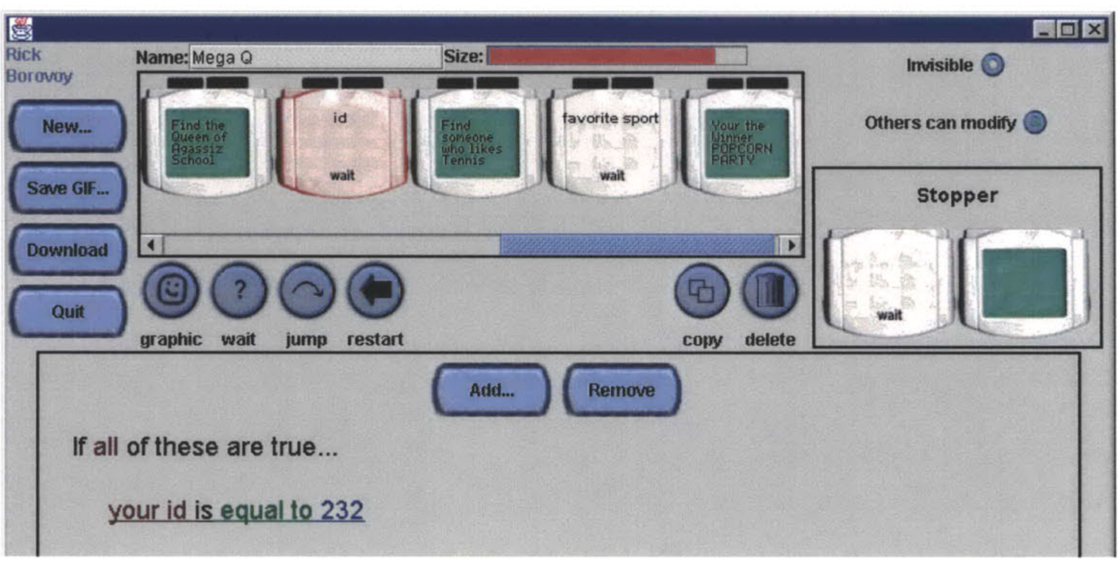

Figure 4-5 Viewing a Rule Block in Mega Quest

Wait rule blocks are straightforward to build, but offer powerful functionality. Using a set of pop up menus, i-ball authors can create rules that test for a variety of conditions relating to personally and socially meaningful dimensions regarding the people running the i-ball and/or the people they are connected to. More of this functionality will be explored in later examples.

Now that some basic functioning of the rule block is clear, we can say more about its implications for building resonant i-balls. I-balls with wait rule blocks can encode their underlying assumptions in much more explicit ways, and require that users enact these assumptions in order to make the i-ball perform. For example, in the case of Mega Q, a quester must discover and accept the author's assumption about whom the queen of the school is if she wants to move on to the next clue. Furthermore, she must enact this assumption by seeking this person out, and physically connecting her device showing the clue to the "queen", further substantiating this belief.

The resonant potential of Mega Q was demonstrated by the strong reaction the gym teacher had when she discovered that she was not the queen of the school. That she wasn't the queen was made vivid for her when some early Mega Q questors tried to connect with her in hopes of advancing their quests, only to leave disappointed. Later, as the word got out that the queen was a sixth grade teacher, students in her midst alerted each other, saying "She's not the queen!" as they ran past the gym teacher in search of the person with the correct id. The gym teacher interpreted this display in terms of the mock rivalry that existed between herself and the sixth grade teacher.

4.2.3 Jump Blocks

Feeling "dissed," the gym teacher created an i-ball in response to Mega Q. "Queen" was designed to start out looking like a regular quest: when students ran the i-ball, an animation instructed them to go see the sixth grade teacher. When they connected to her, however, the i-ball did not offer up another clue, as one would expect with a quest. Instead, it "jumped" from the students' i-socket to the sixth grade teacher's, and displayed a message on her i-socket that said "STOP DREAMING. I AM THE REAL QUEEN IN THIS BUILDING. GUESS WHO"

Figure 4-6 shows the program for Queen. After the "Start" block that begins all programs, there is an animation block, followed by a wait rule block. Like in Mega Q, the wait rule is set up to wait until the user connects to a particular person. In the case of Queen, instead of following the wait rule with another animation that gives a clue, the gym teacher inserted a "Jump" block, that causes the i-ball to copy itself over to an i-socket connected to its "host" i-socket, remove itself from the original host, and then resume execution on the new host i-socket, starting with the program block right after the Jump block. In the case of the Queen i-ball, the next program block is an animation that displays the "STOP DREAMING…" message.

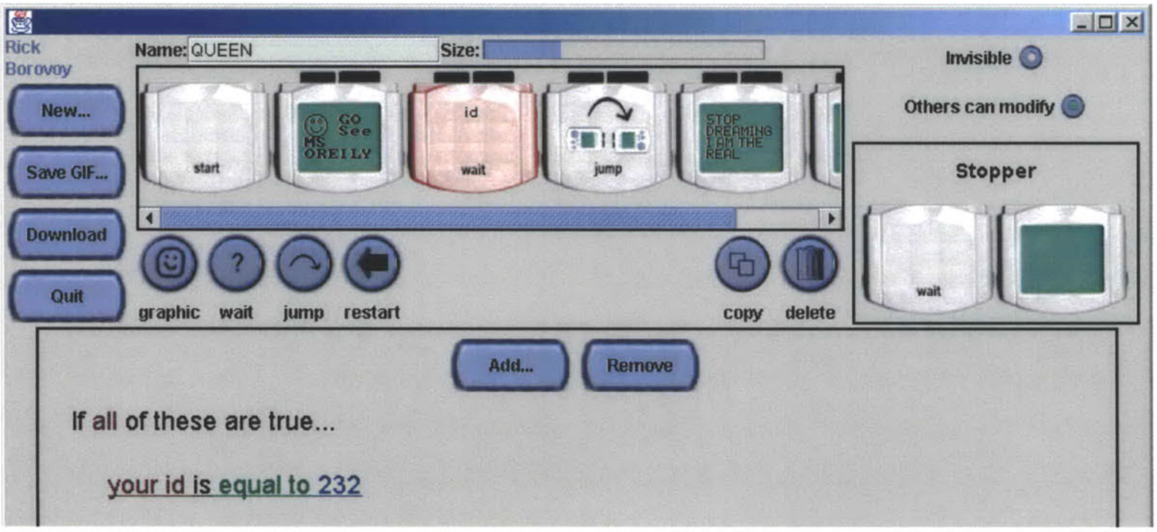

Figure 4-6 The Queen I-ball

Using the Jump and Wait Rule blocks, the gym teacher was able to construct an i-ball that attempted to trick the sixth grade teacher into accepting the "STOP DREAMING" message onto her i-socket. Indeed, the data logs show the trick worked. The sixth grade teacher did in fact wind up with a copy of the Queen i-ball. With the programming environment, the gym teacher was able to author the i-ball's "subtext" so that it matched the surface text: "I'm powerful and clever enough to sneak this message onto your device so I'm the real queen." From the perspective of the gym teacher, tricking the sixth grade teacher into accepting the Queen i-ball was the same as tricking her in to accepting that subtext.

Of course, the more complex the set of beliefs underlying a particular i-ball, the smaller the set of people that will likely resonate with them. The gym teacher created Queen to resonate with her own beliefs and desires, as well as those of her "rival," the sixth grade teacher. The circulation data for Queen reflects this. In Figure 4-7, the top-most box corresponds to the author of Queen. She gave copies of it to the nine people represented by boxes drawn on the row beneath her, and then the i-ball dies. The fact there is very little circulation that does not involve the author suggests that this i-ball did not resonate with a broader audience.

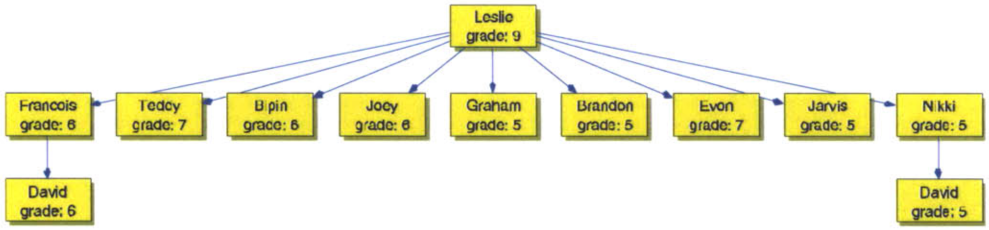

Figure 4-7 Circulation of Queen I-Ball[13]

4.2.4 Mutating Toward Resonance

To make it easier for a wide range of people to author meaningful i-balls, we made it so they did not have to start each one from scratch. Instead, users could open up any i-ball they currently had on their i-socket, edit its contents, and then save it as their own. The adaptable nature of the i-ball programs is one of the hallmarks of folklore, where it serves a similar purpose: to allow those who might not want to or be able to author a new text on their own to rework an existing text to better suit their needs. For a folklore example, consider Henry Jenkins' discussion of "filk" songs: folk-like songs built on the tunes of popular songs by fans of particular television shows. Using this genre, television fans can express and share their thoughts and feelings about TV shows through song without having to compose their own music. We wanted to i-ball programs to have the same adaptive nature as other folk texts, which "are constantly being rewritten, parodied, and amended in order to better facilitate the cultural interests of [their] community"(Jenkins 1992). We also drew inspiration from Amy Bruckman's use of "programming via adaptation" in her work on online virtual communities (Bruckman 1998).

[13] Due to a software bug, this graph does not show that this i-ball actually made it to the "Queen."

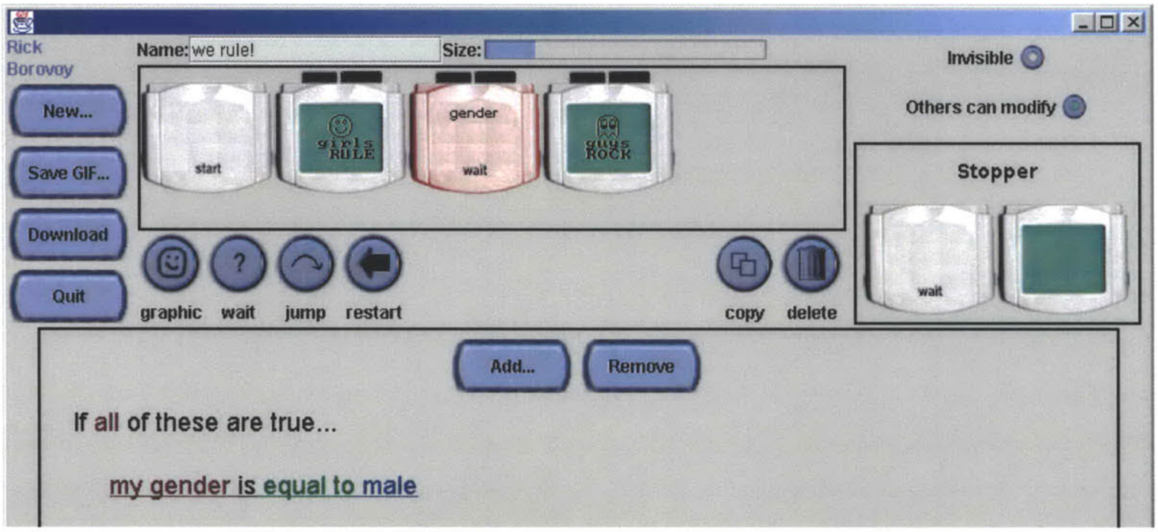

Figure 4-8 "We Rule!" I-Ball

In order to explore how participants adapted other i-balls to reflect their own beliefs and interests, consider the "We Rule!" i-ball that was built by a member of the i-ball research team. "We Rule!" appears to behave like a simple animation: run it and an animation plays. However, it actually shows a different animation depending on the gender of the person playing it: "girls RULE" if played by a female, "guys ROCK" if played by a male. We thought the students at the school might find the functionality of this i-ball compelling. Since the program was a little complex, however, we seeded the user population with an example.

The program for "We Rule!" is shown in Figure 4-8. It consists of two animation blocks separated by a rule block. The first animation block will display "girls RULE". Then the wait rule block executes. It will wait until "my gender is equal to male" before continuing execution. Notice the use of the "my" pronoun here versus the use of "your" in the previous quest example. I-ball programmers choose between "my" and "your" to specify from which i-socket the data gets read. "Your gender" would return the gender of the person whose i-socket is connected to the i-socket with the running i-ball. "My gender" means the gender of the person currently playing this i-ball. Therefore, the i-ball program will pause indefinitely at the wait rule block if the person playing it is a female, and keep going if the person is a male. If the user is a female and execution pauses at the wait rule block, the first animation block will continue to play. However, if the user is male and execution does not pause, the "guys ROCK" animation will play. Here's the tricky part: if a male runs the i-ball, he will only see the "guys ROCK" animation. Although the "girls RULE" animation gets executed first, the program does not pause at the wait rule block, so the "guys ROCK" animation gets executed immediately after. Animation blocks are not guaranteed to play through all of their frames. They play only until another animation block is played. The result is the first animation gets virtually no display time, and so it is invisible.

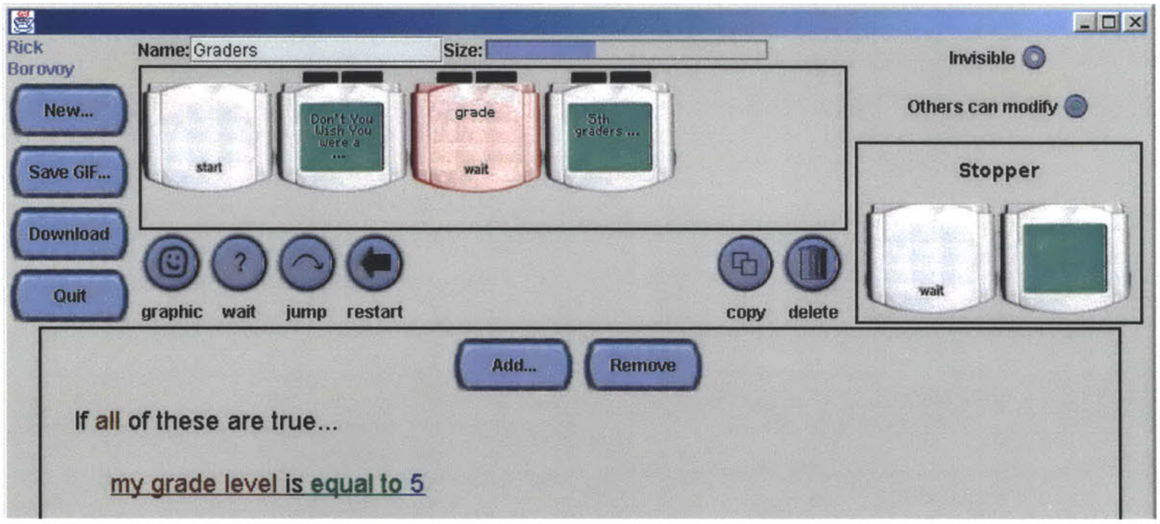

Figure 4-9 The "Graders" I-Ball

The design decision that i-ball animations would run only until the next animation was a complicated one. This behavior was counterintuitive, as evidenced by the number of i-balls made early on in the conference and school trials where users put multiple animation blocks in sequence and expected them to play in sequence (in fact, only the last one plays). Over time, the number of such i-balls diminished, however, as people got comfortable with the "non-blocking" nature of animations. There was one major advantage of non-blocking animations that we felt outweighed their counterintuitive nature: they allowed i-balls to choose between various animations without requiring an explicit branching syntax, enabling a simple i-ball structure that was linear and single-threaded.

It is worth exploring a few of the 78 variations of "We Rule!" that were authored by 40 different participants, to see how people adapted it to their own interests. The original "We Rule!" made somewhat bland, positive statements about the gender of the person running the i-ball. However, many "We Rule!" adaptations reflected a strong belief in the superiority of a group of which the author was a member over another group she was not. For example, the "Graders" I-Ball, created by a fifth grade female (see Figure 4-9), showed an animation that said "5th graders rule, of course" when a fifth grader ran it. When someone who wasn't in the fifth grade looked at the i-ball, it said "Don't you wish you were a fifth grader?" This type of i-ball is similar to a type of humor that works to build in-group solidarity by putting down an out-group and/or celebrating the in-group (Martineau 1972). One key difference is that in those jokes, the insult to the out-group is addressed to the in-group. With the "Graders" i-ball, the insult can be presented as being addressed to the out-group, even as the i-ball circulates within the in-group. This "imagined" delivery of the insult to the out-group might make the i-ball resonate even more strongly with the in-group. Figure 4-10 shows the circulation of the "Graders" i-ball, colorized by grade. Note that it circulated mostly among fifth graders, with the exception of one researcher and a third and sixth grader.

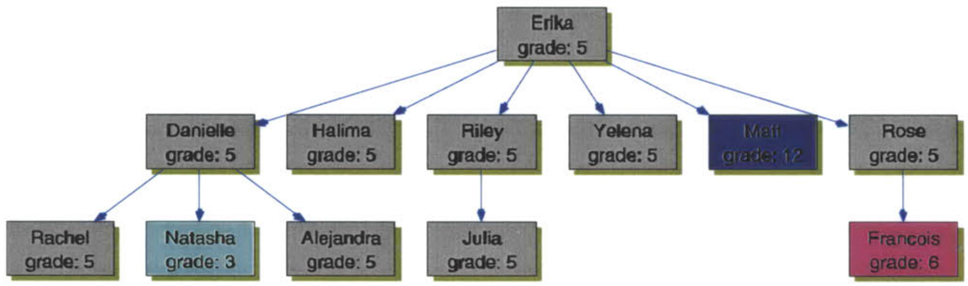

Figure 4-10 Circulation of The "Graders" I-Ball

Another variant of the "We Rule!" i-ball called "Research" (authored by another fifth grade female) uses a different mechanism to resonate with its in-group (see Figure 4-11). When a member of the i-ball research team views this i-ball, it says "hello researchers." When a non-researcher views it, it says, "researchers are weird." In this case, a large part of the appeal of the i-ball is its tacit characterization of the out community as being capable of being deceived into accepting and circulating a text that, while innocuous on the surface, is in fact offensive to them. This negative characterization, articulated through the programmatic structure of the i-ball, is more powerful than the simple assertion that "researchers are weird."

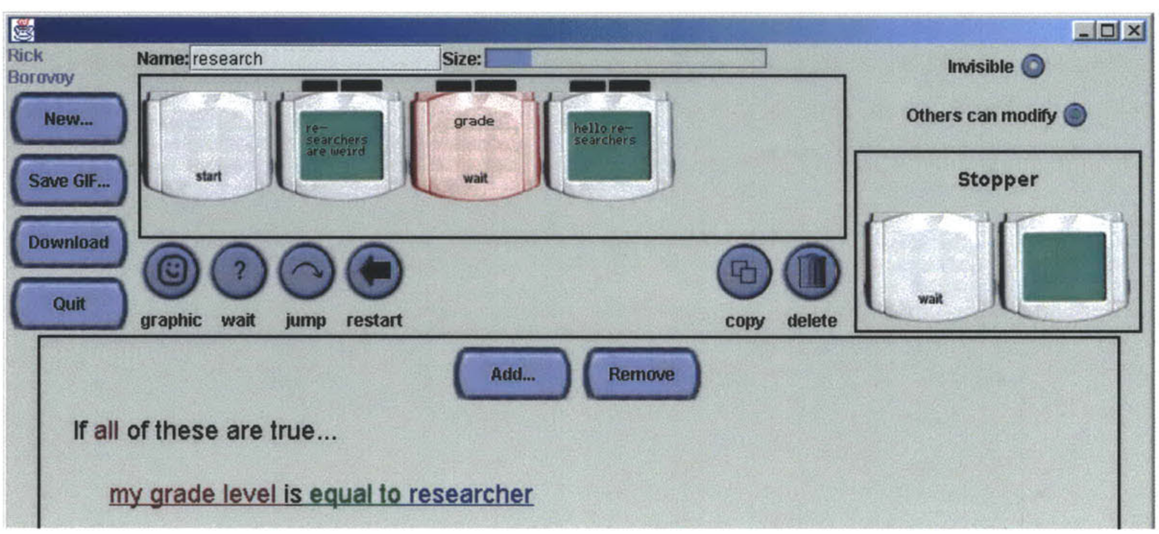

Figure 4-11 The "Research" I-Ball

4.3 Supporting the Oral Circulation of Digital Folk Objects

The i-balls endeavored to increase the expressive potential of the Folk Computing texts, while still allowing them to circulate via informal, face-to-face interaction. In order to accomplish this, we needed a device with several qualities: it had to be wearable, expressive, and readily connectable. In the end, we found an off-the-shelf device manufactured by SEGA that met our hardware needs, and we developed our own software for it.

All Folk Computing devices to date have been wearable. Early on with the Thinking Tags, we discovered that wearability was one way to reduce the barriers to participating in the folklore process. Wearing a device identified oneself as a participant and served as an invitation for others to approach and initiate a transaction, thereby reducing the social risk of such an overture. Therefore, as we started to explore more complex and expressive devices, we still wanted something that could be visible displayed on one's person.

A wearable device was also more likely to accompany its owner into the "off the grid" locations in which folklore exchange occurs. Verbal folklore is pervasive because no matter where one goes, one's mouth comes along. We wanted the i-sockets to function as a new kind of mouth, capable of different kind of expression, but always available like any other organ.

In order for a device to be wearable, it had to be small and lightweight, as well as durable and at least modestly fashionable. Fortunately, the SEGA device was designed for the children's market with lightness and durability in mind. In terms of fashion, however, the device seemed a little dull. Therefore, we spent considerable time looking for attractive ways for people to suspend the devices around their necks. In the end, we offered kids at the school a choice between multicolored yoyo strings and ball-chain – the type used in the military for ID tag necklaces. The ball-chain became the "cool" choice, partly because it made fun sounds and enabled various i-socket spinning activities (see Figure 4-12)

Figure 4-12 A Playground Demonstration of I-Socket Spinning

Although the graphic capabilities of the i-socket – 48 by 32 monochromatic pixels – are modest by PC, or even PDA standards, they provided adequate expressive range for the i-ball trial for several reasons. First, kids were able to author their own animations and games for this display, which gave them a different relationship to the content. Second, there were no examples of professionally designed content for the i-sockets to which participants had to compare their work (this is not the case when kids design their own video games for a PC). Third, when people see the small, monochrome LCD display, they scale down their expectations about what it is possible to produce. Finally – like the similarly small PostIt Note – the expressive power of the i-socket screen was amplified by its ability to venture in to meaningful contexts. When all these factors combined, people were excited to be able to make a digital artifact that they could share with their friends during the times and in the places that were most meaningful to them.

The final requirement for the i-socket was connectivity. Here, we were aided by the fact that SEGA had designed their devices to easily connect to one another, or to dock with a base station. We added software that made it trivial for a users to copy i-balls between i-sockets. Indeed, many users figured out how to do this without instruction immediately after receiving their i-sockets. In this way, users were able to make arbitrarily complex i-balls, and then seamlessly make copies of them for as many people as they wanted, wherever they wanted.

4.4 Building Visualizations of Patterns of Resonance

To complement the tools for authoring, performing and exchanging resonant objects, we needed a visualization tool to let participants explore the circulation patterns of these objects. In this way, users could create personally meaningful objects, release them into the community, and follow their movements as they traced the edges of like-minded subcultures.

While the "community mirror" visualizations for the Meme tags were directed toward a large audience, we designed the i-ball visualization tool to be used by individual participants. The Community Mirrors were designed to display broad-based community trends in the background. For the next generation, we designed an interactive tool that would allow people to inquire more deeply into the circulation patterns of i-balls.

We wanted participants to feel a personal connection to the visualizations. Children understand visualizations better when they can locate familiar objects in them (Hancock, Kaput et al. 1992). To that end, we designed most of the visualizations to allow people to explore data about them as individuals and their relationship to the larger population. For example, rather than potentially drowning participants in the data of several thousand i-balls, we let them explore data only about the i-balls they had held on their i-socket at some point. From an interface perspective, this made for a much more manageable list of i-balls to choose from. From a learning perspective, these i-balls turned out to have tremendous leverage. Based on participants' familiarity with these i-balls – e.g. who they got them from, who they gave them to – they already had preconceptions about the i-balls' circulation patterns. We observed many occasions where the discrepancy between these preconceptions and the patterns of circulation revealed by the visualizations created a powerful learning moment.

In order to make the visualizations feel relevant, we wanted them to contain data about people and/or i-balls with which the viewer was familiar. The fundamental unit in almost all i-ball visualizations was the "i-ball reception." An i-ball reception occurs when a particular person receives a copy of a particular i-ball. We experimented with a variety of i-ball reception visualizations, each of which falls within a two dimensional space according to how data is aggregated (see Figure 4-13). Some visualizations, such as the top ten list of a particular author's most popular i-balls, aggregated i-ball reception data across participants. These visualizations showed data about particular i-balls but not about particular people. Other visualizations, such as the ego network visualization of a particular person, did the opposite. They aggregated data across i-balls, but showed information about particular people.

	Individual I-Balls	Aggregated I-Balls
Aggregated Participants	I-Ball Popularity Top Ten Lists Individual I-Ball Demographics	I-Ball Author Demographics
Individual Participants	I-Ball Travel Visualizaitons I-Ball Mutation Visualizaitons	Social Network Visualizations

Figure 4-13 Design Space of I-Ball Visualizations

The following sections describe the visualizations we experimented with, as well some ideas for new visualizations based on what we learned. There is a chicken-and-egg problem with designing visualizations for a large-scale Folk Computing experiment: it is hard to design meaningful visualizations without a corpus of data to experiment with, but you can not collect that corpus without having run a trial that includes some visualizations. Now that we have collected a rich set of i-ball interaction data, we have the opportunity to experiment with new visualizations that could be used in a future trial.

4.4.1 I-Ball Specific Visualizations

Based on our experience with the Meme Tag "Mindshare" community mirror, we wanted people who had authored or carried a particular i-ball to get information about its resonant population, broken down by key demographic dimensions (e.g. grade, gender, ethnicity). Figure 4-15 shows two graphs that represent the set of people who received the "B & F" i-ball – an "advertisement" for a book that two sixth grade boys were working on (see Figure 4-14). The pie chart of the left shows that of the 42 people who got copies of this i-ball, 81% (34) were male and 19% (8) were female. The bar chart on the right shows that the i-ball had carriers in five different grades, with sixth grade having by far the largest number of carriers.

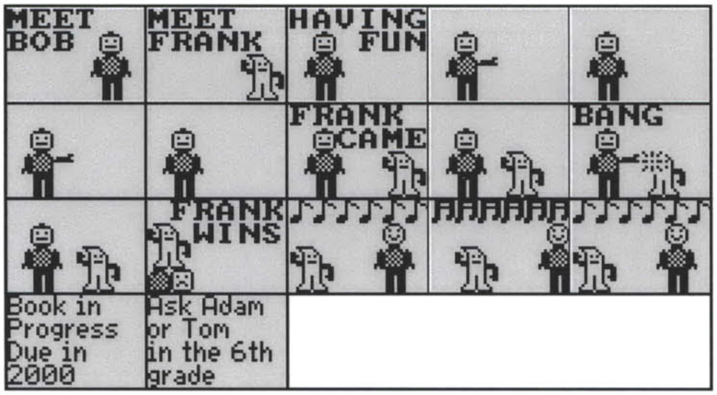

Figure 4-14 "B & F" I-Ball Animation Frames

These i-ball demographic graphs were most popular with teachers who wanted to use them to help students meet the "data literacy" requirements of the Massachusetts MCAS standardized tests. A few teachers facilitated data analysis and discussion sessions, and kids were able to use these graphs to explore issues such as gender bias in particular i-balls.

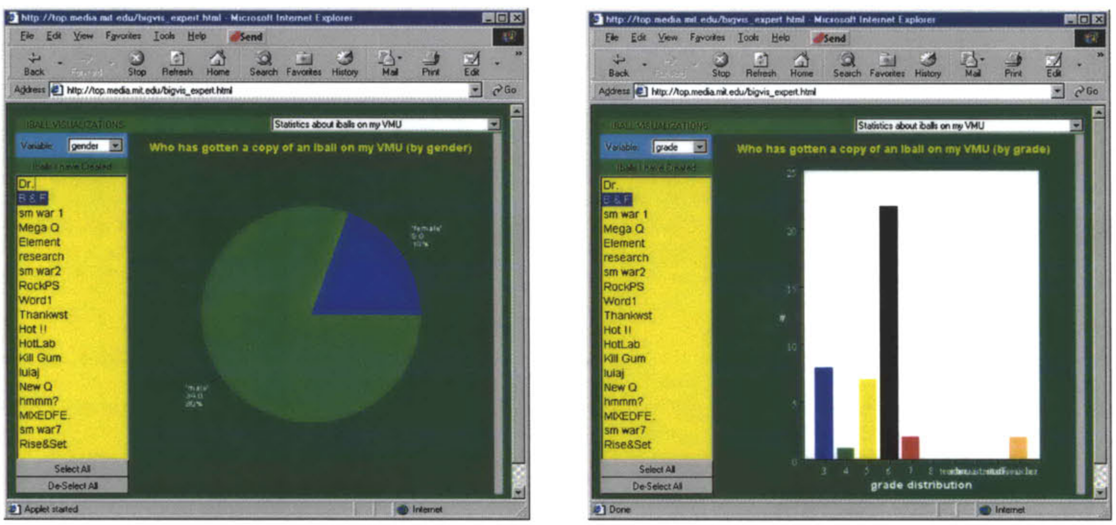

Figure 4-15 Visualizations of the Audience for a Particular I-Ball

Many people suggested that it would be useful to have data about how an i-ball circulated available alongside the i-ball, on the i-socket device. If circulation data could be delivered in situ, it would be more accessible and meaningful to the participants, and could be taken up more easily in conversation. Such visualizations would be more like the "routing slips" that provide readers with circulation information for a variety of documents, including the "borrowing" card in library books and the cc field in email messages. Such routing slips help those who read them understand something about the significance of the document to the community. As Sealy Brown and Duguid suggest:

> We recognize important organizational documents from the long and impressive routing slips attached... And we assert our membership in a community in part by showing we have read those documents -- which is why we often like to be sure our own name gets on the routing slips... (Brown and Duguid 1996)

The challenge is to design a "routing slip" for i-balls that displays key data about the i-balls' circulation within the constraints of a small, portable device like the i-socket. The i-ball "saturation slips" shown in Figure 4-16 are one possible solution. These slips use a grayscale value to represent how "saturated" a particular demographic group is with a particular i-ball. Saturation values are given as percentages, where 100% -- represented by a solid black bar – means that everyone in that group got a copy of the i-ball, and 0% -- represented by a white bar – means no one got it. In the case of the "B & F" i-ball, one can quickly determine that the circulation of the i-ball was skewed significantly toward 6th grade and toward males, although there was some circulation among several grades and both genders.

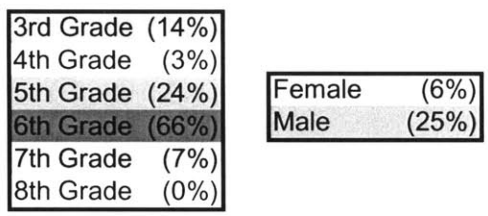

Figure 4-16 Saturation Slips for "B & F" I-Ball

The "My Top Ten Most Popular I-Balls" visualization gave participants a basic measure of how widely the i-balls they had created had circulated (see Figure 4-17). We also built a Top Ten List for the most popular i-balls of all time. As with the Meme Tags, this turned out to be one of the most popular visualizations, as there was a lot of status associated with getting an i-ball on this list.

Figure 4-17 Top Ten Most Popular I-Balls Visualizations

The popularity of the "Top Ten List" form as a means for getting a quick read on community taste could also be applied to sub-populations. Figure 4-18 shows two "comparative top ten lists" that tell which i-balls were most popular with each gender. The yellow highlighting shows "signature i-balls": i-balls that are unique to that gender's top ten list. With this visualization, one can quickly determine that "hquest", "Mega Q", "Romance" and "RockPS" were popular with both sexes, whereas "boom", "WWII", "Gumby", "Mr. Ghost", "Darat", and "Mr. BOMB" were more popular with males. Interestingly, with the exception of "Darat", each of these i-balls is somewhat violent, whereas none of the i-balls that were uniquely popular among females have any violent content. Although gender skew in the appeal of violent content is somewhat stereotypical, it suggests that this simple visualization has the potential to reveal interesting patterns of shared and differential resonance.

I-Ball	# Males	I-Ball	# Females
hquest	111	hquest	95
Mega Q	74	romance	70
romance	61	Mega Q	66
boom	56	Agassiz	41
WWII	55	RockPS	31
Gumby	46	Rashaad	31
Mr.Ghost	43	we rule!	30
Darat	43	Shadow	30
RockPS	39	Boston	29
MR.BOMB	39	Mega Q3	28

Figure 4-18 Comparative Top Ten Lists by Gender, with "Signature" I-Balls in Yellow

The goal of the "Comparative Top Ten list" is to allow people who resonate with particular i-balls to learn about who else resonates with them, and therefore who else might share some of the underlying knowledge, experiences, beliefs, and desires associated with them. On the flip side, they also learn about the boundaries of these resonant groups. For another example, we can see what a comparative top ten list by grade level could reveal to members of the school population about the distinctive commitments that unite and divide people in different grades (see Figure 4-19) Consider some of the signature i-balls of various grades. "Shadow", which is popular only in the third grade, depicts an eponymous third grade class pet rabbit. "O'Reilly", which is uniquely popular in the sixth grade, is actually a homework assignment handed out in i-ball form by the sixth grade teacher of that name. Finally, "Rashaad" is an animation that shows a cityscape and says "Cash! Money! Dice! The block is hot." It was uniquely popular in the eighth grade, which is around the age that many kids get interested in gambling (http://www.teengambler.com/facts/facts.html) Through this type of visualization, the i-ball "in group" represented by each of these examples has the opportunity to learn something about the narrowness and breadth of its constituency.

3rd	4th	5th	6th	7th	8th
hquest	hquest	hquest	hquest	hquest	boom
romance	romance	Mega Q	Mega Q	Gumby	Rashaad
Shadow	Agassiz	Mr. Fuzz	O'Reilly	romance	bobby
Mr. Fuzz	Darat	WWII	highland	Mega Q	hquest
WWII	lilmac	we rule!	B & F	s.kill	Mega Q
Mega Q	Smilely	blows up	O'Reilly	boom	s.kill
Agassiz	robot	RockPS	Hyland	we rule!	Agassiz
RockPS	Mega Q3	Mr.Ghost	we rule!	death	romance
Dr.	Chris	lighting	Catch	Boston	Hello
robot	airman	Darat	DEATH!	HERCULE	hquest

Figure 4-19 Comparative Top Ten List by Grade

4.4.2 Visualizations that Reflect Individual I-Ball Exchanges

In order to help people better understand the social structure within which the i-balls traveled, we wanted to create some visualizations that would reflect the dynamics of i-ball circulation. While i-ball demographic visualizations like the ones in Figure 4-15 and Figure 4-16 provide information about who got a copy of a particular i-ball, they do not show the process by which these i-balls spread through the population. We believed that visualizing the process of an i-ball flowing through a population would help people overcome their ego-network-centered understanding of their social network.

We chose to focus on visualizing the circulation of single i-balls, rather than combining the data from multiple i-balls into a traditional social network diagram. Figure 4-20 shows a social network diagram on the left, and an i-ball network visualization on the right. The top-most node of an i-ball network represents the i-ball's creator. Nodes on the level below this i-ball and connected to it by a blue line represent the people who got the i-ball from the parent node, and so on down the tree. Individual nodes are labeled with the name and grade level of the person they represent.

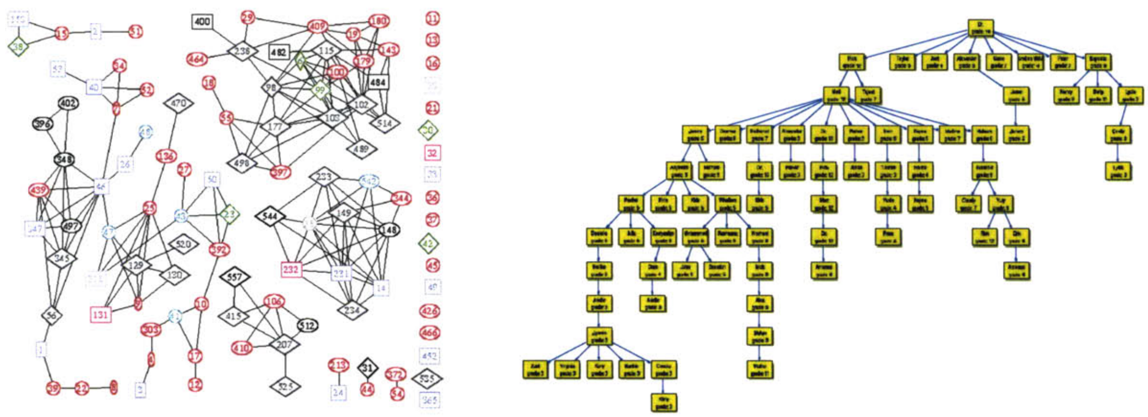

Figure 4-20 Social Network Diagram (Left) and I-Ball Network Diagram (Right)

Of course, i-ball networks are limited in comparison to social networks in terms of what they reveal about the structure of a community. An i-ball network is a snapshot of interactions within a community that occurred around a particular i-ball. A social network visualization could composite data from many i-balls in an effort to reveal more persistent patterns of interaction, such as cliques. Other research tools have produced social network visualizations by analyzing interaction data (usually using online sources such as email or Usenet logs) (Sack 2000). While it would certainly be useful to apply some of those tools to i-ball data, we chose to focus on exploring the unique opportunities for helping people understand community dynamics offered by the ability to track the copies of a single object as it spreads through a population.

I-ball network visualizations were designed to help participants reconcile their experience of an i-ball – whom they got it from and whom they gave it to – with a community-wide perspective that is usually invisible. Research in children's folklore has shown that children have an ego-network-centered perspective on folklore: they often think that the person who shared it with them invented it, and that the people they gave it to are the last people to get it. By locating themselves on an i-ball network diagram and confirming that it included their experience with the i-ball, users were then able to enlarge their understanding of the i-ball's behavior by tracing the links that went beyond their direct experience.

There are several reasons why an i-ball network visualization may be superior to a social network visualization for helping people begin to move beyond an ego-network-centric understanding of community structure. First, because of the simplicity and uniformity of the relationships represented by the i-ball network, it is possible for someone whose interaction comprises a small piece of the network to identify with the rest of it. A link between two nodes on an i-ball network means that one person passed another person the particular i-ball represented, so it is easy for a user to imagine that the links that do not involve her are similar to the ones that do. This is less so with social networks, whose links represent one or many highly varied interactions between people.

The implicit temporal organization of an i-ball network is also important to people's ability to conceive of its totality. I-ball networks flow from top to bottom, representing multiple generations of i-ball transmission. That makes it possible for a user to conceive of the i-ball network as a single super-event unfolding over time in which he participated. He can locate himself in this event not only in terms of other participants, but also in terms of time. Such a temporal dimension is missing in social networks.

In future work, we would like to combine the advantages of the i-ball network and the i-ball demographic visualizations. Pie and bar charts like in Figure 4-15 show overall demographic data for an i-ball audience, but compress out all information about how the audience grew over time. I-ball networks show how i-balls spread through pair-wise interaction, and were more compelling to users, but they do not do a good job of highlighting patterns relating to the characteristics of participants.

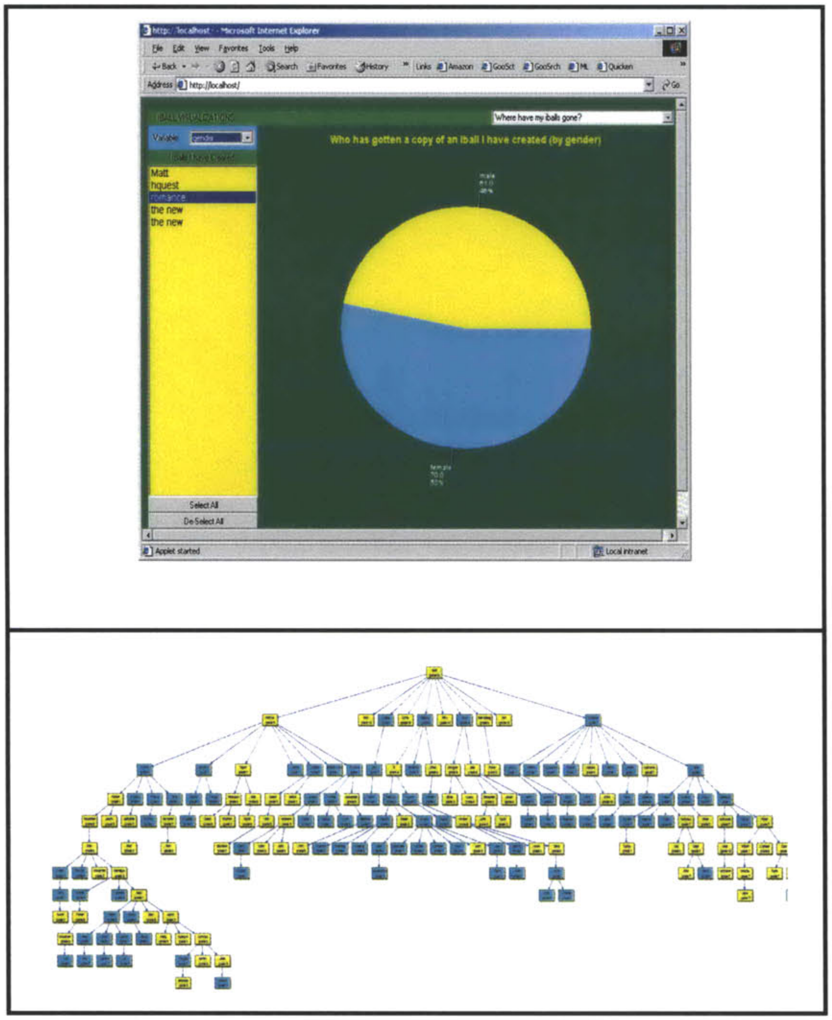

Figure 4-21 Visualizing Gender of Romance Audience:
Pie Chart (Top) Vs. Colorized I-Ball Network (Bottom)
(Female is Yellow, Male is Blue)

The bottom part of Figure 4-21 shows a new type of visualization designed after the school trial to add a demographic layer to the i-ball network visualization. Comparing this visualization to the pie chart above it shows how much more information the colorized i-ball network reveals. While both graphs show that roughly equal numbers of males and females got the Romance i-ball, the colorized network reveals a gender interaction pattern that is otherwise hidden: namely that participants had a tendency to give Romance to others of the same gender, resulting in clusters of blue and yellow nodes in the network diagram.

Apart from the colorizing, the i-ball network visualizations had other features that made them more compelling than the simple pie and bar charts. One such feature was the network diagram's unique ability to represent the magnitude of the i-ball's popularity in a way that is absolute, not relative to a scale or specified by a label. The larger the i-ball's audience, the more nodes appeared on the graph. Contrast this with the pie and bar charts, whose shapes – without further annotation – reveal nothing about the size of the population. The shapes of the i-ball networks are also more interesting and more varied than the shapes that comprise the simpler charts. For example, the i-ball network visualizations in both Figure 4-20 and Figure 4-21 both contain protruding sub-graphs that are each spawned by single nodes, reflecting something potentially interesting about those people.

Having discussed some of the features of I-balls, we will next turn to practical applications of the underlying Folk Computing paradigm in educational settings.

5 Folk Computing and Education: Supporting an Ecology of Learners

Much computer-supported collaborative learning (CSCL) research has focused on innovations designed to work either within the traditional scope of a classroom, to bridge several remotely located classrooms, or to work entirely outside of the school structure (Koschmann 1996). We attempted to integrate i-ball technology into the school's entire ecology, defined as "the totality or pattern of relations between organisms and their environment" (Merriam-Webster, 2001, second def.). Our aim is to use technology to create an "ecology of learners" that moves beyond the temporal, spatial, and social boundaries of an individual classroom by taking better advantage of the ecology of the entire school, including: the various populations of students, teachers, and staff that inhabit a school; the social network of relationships that exist within and between these groups; the socio-temporo-spatial "niches" that represent various individual classes; and a wide range of settings with less formal structures, such as the hallways, the stairwells, the lunch room, the playground and the school bus.

In a technology-supported "ecology of learners", there is a mutually beneficial relationship between the technology and school ecology. The technology aims to integrate with, probe, and enrich the ecology. The ecology then becomes a meaningful background against which people can engage with the technology. This has two important learning outcomes. First, the school ecology itself becomes an object of scientific inquiry, with students playing a dual role of both ecologists and ecosystem inhabitants. Second, this inquiry, as well as inquiries into other subject domains, generates additional meaning and enjoyment from its connection to a rich ecological context. Our goal is to extend the "community of learners" (Rogoff 1994) concept beyond classmates interacting in a classroom to include siblings having a learning exchange in the hall and friends in a range of grades sharing some knowledge over lunch.

Simply providing a school with ubiquitous computing technology is not enough, however. In order to support a true ecology of learners, the technology must be built around a model that meets several requirements. First, the technology must be friendly to the school ecosystem: it cannot inherently place too many demands on who can interact where or when, thereby imposing its own structure on top of the pre-existing ecology. Second, it must transparently collect data about the school ecology, which can then become an object of study. Third, it must be easy and customizable enough that it will be taken up by teachers and students who wish to appropriate it for their "ecologically relevant" activities in their own ecological niches.

Supporting a school-wide ecology of learners leads to two main educational outcomes: new opportunities for students to inquire into the complex system that is their own school ecology, and new ways of making a variety of educational goals more meaningful and engaging to students by aligning them with this ecology. The two main parts of this chapter will explore how Folk Computing technology can be used to achieve each these outcomes.

5.1 New opportunities for studying complex systems

We want students to develop deeper understandings of ecological systems, and of the ecological paradigm itself (Resnick 2001). To this end, our technology lets students play a dual role as both a participant in an ecosystem and as an ecologist studying that system. This theme has been explored in prior research on Participatory Simulations (Colella, Borovoy et al. 1998). Our current research moves beyond simulation by allowing students to study their own real, pre-existing ecosystem by constructing their own "probes", releasing them into the ecology, and then tracking the whereabouts of these probes as they move through the system. In the following sections, we will first establish that a school-wide ecology existed, and that the i-ball activity did in fact integrate with it. We will then explore how the i-balls that students created acted like scientific probes with which they could explore this ecology.

5.1.1 I-Ball Integration with School Ecology

A typical Participatory Simulation event involved 20 or 30 participants trying to figure out how a simulated virus was spreading amongst themselves by conducting several experiments over the course of a few hours. This was a very innovative activity in terms of how it made scientific inquiry feel personal and social, intimate and analytical. From the perspective of supporting a learning ecology, however, the activity was fairly traditional: it brought a relatively narrow selection of people into a constrained space for a short period of time. Other research into this type of activity exhibits the same constraints (Danesh, Inkpen et al. 2001; Soloway, Norris et al. 2001).

The i-ball activity was designed to put fewer constraints on where, when, and what interactions between participants were supposed to take place. Our goal was to preserve the students' school ecology, so the school ecology itself – rather than a simulated virus ecosystem – could become an object of study. For a particular participant, we wanted to ensure that the activity included not just students in her class, but other school friends, other siblings at the school, and other teachers and staff she might have contact with (including the principal and lunch-room monitor). We also wanted to make sure the activity could work before, during and after school in the whole variety of school settings. The i-ball activity is modeled on the exchange of children's folklore, partly because children's folklore is known to be "eco-friendly": that is, performance and exchange of jokes, legends, games, and rhymes is already a part of many of the interactions between school children that take place in a variety of school contexts.

	3	4	5	6	7	8
3	973	146	183	142	78	62
4	146	457	61	124	31	40
5	183	61	743	103	28	26
6	142	124	103	860	103	64
7	78	31	28	103	444	106
8	62	40	26	64	106	604

Figure 5-1 Number of I-balls Exchanged Between Grades 3 through 8

Figure 5-1 shows the number of i-balls that were exchanged within and between grades 3 through 8. In order to make the pattern of interaction more clear, each cell has been "colorized" according to how much exchange happened between those classes: the darkest cell (third graders exchanging amongst themselves) had the most activity, the lightest cells had the least. The dark diagonal "line" from the upper-left to lower-right of the table shows the tendencies for people in the same grade to exchange i-balls with each other. This reflects the significant role that grade level plays in organizing social interaction in a school. However, it is clear that what happens within a particular grade is only a part of a larger pattern. Note that each grade has some interaction with every other. One can also see the somewhat faint images of two schools within the school, with third through fifth grades in one sub-ecosystem, and seventh and eighth grade in another, with little interaction between the two. Sixth grade provides a bridge between these.

Factor	Odds Ratio	Significance
Same last name	21.60	P<.0001
Same classroom	6.21	P<.0001
Same grade	4.95	P<.0001
Same gender	2.53	P<.0001
Same floor	1.71	P<.0001

Figure 5-2 Analysis of Factors That Affected the Probability of an Interaction Between Two People

It turns out that whether or not two students are in the same grade is not the strongest predictor of whether they will exchange i-balls. Figure 5-2 shows the results of a logistic regression, which reveals factors that had the greatest role in determining whether a randomly chosen pair of students in the school population exchanged an i-ball ($R^2 = .29$). If the classroom unit were the only organizing structure that influenced who interacted with whom, there would be no other statistically significant factors. This is not the case, however. The strongest factor was "same last name", which should be considered a first-order approximation of "same family". Comparing odds ratios reveals that sharing a last name had more than three times more impact on whether two people interacted than "same class."

Siblings are an excellent example of a type of relationship that plays an important role in a school's ecology, but that is mostly excluded from its formal structure. A lot of i-ball activity became additionally meaningful because of its ability to participate in these relationships. Three are several examples of exchanges like the one that involved two brothers playing "catch" with a hitcher i-ball at 10pm on a Monday night.

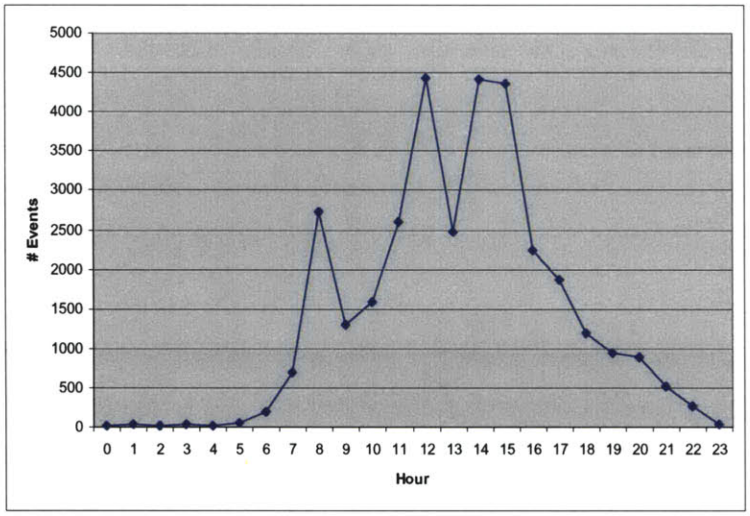

Figure 5-3 Total I-Socket Activity by Hour of the Day

Figure 5-3 shows the amount of i-socket activity that took place during each hour of the day, summed across all users during each day of the trial. The three peaks correspond roughly to the period right before school starts, lunch, and just after school. Notice that i-ball activity after school trails off quite slowly, with more than 10% of i-ball activity occurring after 5PM. 32% of i-ball activity occurred outside of school hours. These data suggest that the i-ball activity found its way into all segments of the school day.

Figure 5-4 shows some of the many settings where people at the school played with and exchanged i-balls: the classroom, the office, the playground and the school bus. With the current i-ball technology, we were not able to track where i-ball exchanges occurred. Therefore, we had to rely on observation and user reports to determine where i-ball activity occurred. These sources confirmed that there was substantial i-ball activity at various times in most school locations, from the teacher's room to the lunchroom to the bathroom.

Figure 5-4 I-balls Usage in a Variety of School Locales

5.1.2 Constructing StarProbes to Investigate School Ecology

Because the i-ball activity permeated the school ecology, i-balls themselves became powerful diagnostic tools for exploring this ecology. This was especially true because students created i-balls themselves. Their investment in those i-balls, combined with their desire to influence their peers, made students very interested in how their i-balls spread through the population.

As described in Chapter 4, we built a suite of visualization tools to allow students to track their i-balls. In the spirit of the Logo turtle, we tried to preserve the "object to think with" status of the i-ball in our visualizations (Papert 1980). Therefore, we focused on visualizations that preserved the object-quality of particular i-balls, rather than creating more traditional and abstract social network diagrams (Wasserman and Faust 1994), derived from the amalgamated data of multiple i-balls. There is support for this type of approach in traditional children's activities. "Message in a bottle" and "Message in a balloon" activities can both be seen as opportunities to construct primitive scientific "probes," where the trajectory of a single element will shed light on the dynamics of the ocean and atmosphere, respectively. These probes are, like the LOGO turtle, body syntonic: you identify with them and share in their experience as they make their way through a complex space. Another example is the popular school activity where a classroom will create a stuffed animal as a surrogate for themselves, equip him with a camera and self-addressed envelopes, and then send him off to another class in hopes that he will travel around the world and "report" on his adventures. Chain letters also have this quality, with the additional property that the probe replicates as it travels.

We have created the term "StarProbe" to describe a probe object that helps reveal the structure of a complex system by reporting on its pattern of circulation within the system. They are inspired by StarLogo, a Logo environment that takes an object-centered approach to modeling complex systems by representing them as thousands of interacting, relatively simple turtles (Resnick 1999).

StarProbes are also part of the adult world. Federal Express packages are StarProbes when one monitors their progress across the country via tracking on the World Wide Web frequently enough to learn about the structure of the FedEx system. The popularity of a site like WheresGeorge.com, which allows people to track dollar bills as they circulate around the U.S., is testament to the appeal of creating and tracking StarProbes. Medical diagnosis frequently makes use of dyes whose patterns of circulation, after being injected into a vein, reveal much about the internal structure and dynamics of the body. The popular movie Twister highlights a set of tennis-ball size probes that are designed to be injected into a live tornado and tracked to reveal its inner structure.

Of course, one of the main differences between the "message in a bottle" StarProbes of children and the "radioactive dye in a body" probes of scientists, is the quality of the feedback. Messages in bottles, chain letters, stuffed animals, and WheresGeorge all rely on human generated feedback about the whereabouts of the probe, which quite often never comes. We carefully constructed the i-balls to function as robust and autonomous StarProbes, tracking their own circulation, and automatically report this data back to a central server.

5.1.3 Examples of Students as Ecologists

Unfortunately, from a technical perspective, the visualization tools turned out to be the least robust part of the i-ball system. The computers at the school were too slow and the screens too small to run many of the visualizations effectively. In an effort to compensate, we brought in a few additional higher-powered machines for students to run visualizations on (see Figure 5-5 Right). We also made large poster-sized printouts of some of the most popular i-balls and brought them to the school (Figure 5-5 Left). Nevertheless, technical challenges kept the visualization component from becoming an integral part of the i-ball experience. This is something we aim to correct in future versions, ideally by allowing students to view visualizations using the same devices used for exchanging i-balls.

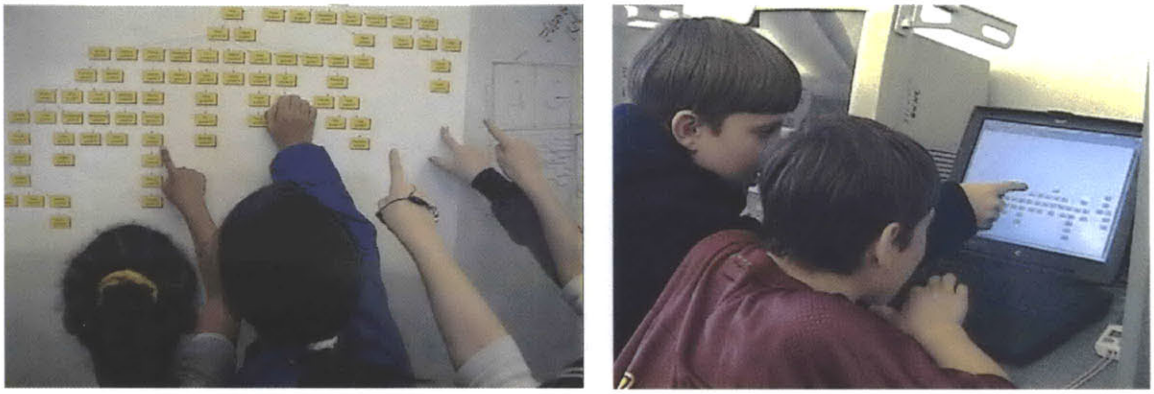

Figure 5-5 Students Tracking I-Balls

As the images in Figure 5-5 reflect, to the extent they were able to, kids got very involved in the process of tracking how their i-balls moved through the school. The tracking visualizations, which look somewhat abstract to outsiders, were very meaningful to them. Whenever we put up poster-sized versions of popular i-balls, students would swarm around looking for their names, exploring things like how early they were in that particular i-ball "trend," and what role they played in further spreading the i-ball. Comparing their role to that of their friends was also a popular activity. Furthermore, it was not just students who got engaged in tracking their i-balls. The computer teacher at the school created one of the most successful i-balls, called "MegaQuest", and he was very excited when we presented him with a large-format visualization showing how it spread. In fact, when we conducted a follow-up interview with him more than a year after the trial, the 10' x 3' poster was still hanging in his room, with a sign he attached that said "Mr. –'s MegaQuest i-ball spread to 141 people."

For many students, the visualizations provided their first glance into the structure of a social realm that existed beyond their direct experience, beyond the "first degree of separation." This has important implications for school-age kids, whose social reality is dominated by the construct of cliques. A clique is defined in terms of direct, single-degree connections. Specifically, it is a set of people within a larger population who each have a direct, single-degree relationship with everyone else in the set (Wasserman and Faust 1994). Because people in cliques still have many connections outside the clique, cliques become less important when second-degree, friend-of-friend relationships are considered significant. The i-balls activity helped many kids experience the power of these relationships for the first time. By watching their i-balls move through the population, users were able to observe how their influence spread beyond the first-degree horizon of their immediate social circle.

For example, one girl, upon viewing a primitive visualization that listed who had received a particular i-ball she made exclaimed, "How did she get a copy of my i-ball? I didn't give it to her. I don't like her." After further analysis, she realized that the i-ball had reached the other girl by way of a friend. This same type of realization occurred when we asked other students if they could tell us exactly who had an i-ball that they made. Several expressed with great confidence that they could, and then were surprised by what they saw in visualization. Often, the i-ball was more popular than they had thought, because people they had given it to then gave it to others, and so on.

One third grade class was surprised when a visualization of one of their favorite i-balls revealed it to be less popular than they had imagined. Someone in their class had created an i-ball likeness of Shadow, their class pet bunny. They were sure this i-ball was very popular throughout the school and were disappointed when they saw it had not spread that far beyond their classroom. Seeing the visualization and discussing its significance caused them to consider that their perceptions may have been influenced by a small sample bias: they experienced the i-ball spreading in their classroom niche, and assumed that was representative of the larger population.

As mentioned previously, misunderstandings about how a piece of knowledge is distributed in one's larger social system are common. Unfortunately, there is usually very little chance for people to correct their misunderstandings about their social environment. As in the above examples, they are often unaware of the limits of their knowledge, and the closed social circles they travel in keep them from encountering anything problematic. The i-ball technology began to provide an opportunity for people to see outside their own local context and get a glimpse of the larger patterns. Because people were interested in how influential their i-balls were, they were motivated to examine the visualizations. Since these patterns were constructed out of the traces of probes they themselves had created, the resulting visualizations were comprehensible to them. Because these patterns were made up of people they knew or knew about, the visualizations bore particular significance.

We end this section with some ideas about how the i-ball technology could enable students to study the diversity of the school ecosystem. The school was very interested in celebrating and reflecting on its racial diversity. During the trial, students were able to examine graphs representing the ethnic make-up of the recipients for a particular i-ball, to see if it reached a diverse audience. The colorized i-ball network visualizations that we developed after the trial seemed particularly effective for exploring the concept of diversity, however. For example, consider the visualization shown in Figure 5-6 that highlights the ethnicity of all the people who received the "Romance" i-ball. This graph shows not only that the i-ball reached an ethnically diverse group, but also that it was frequently passed between people of different ethnicity. This is not the case in Figure 5-7, which is colorizes the same network based on gender (we

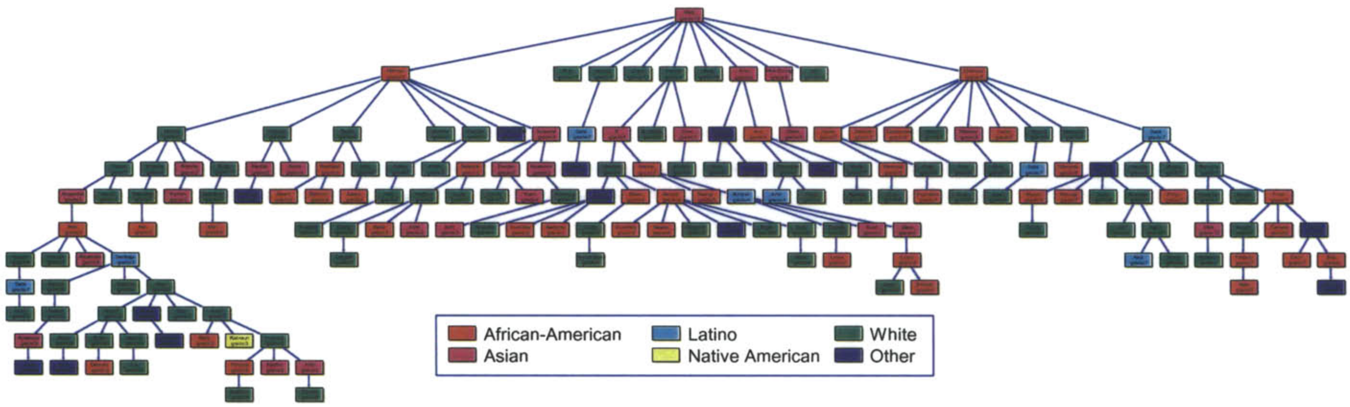

Figure 5-6 Network Visualization of "Romance" I-Ball, Colorized by Ethnicity

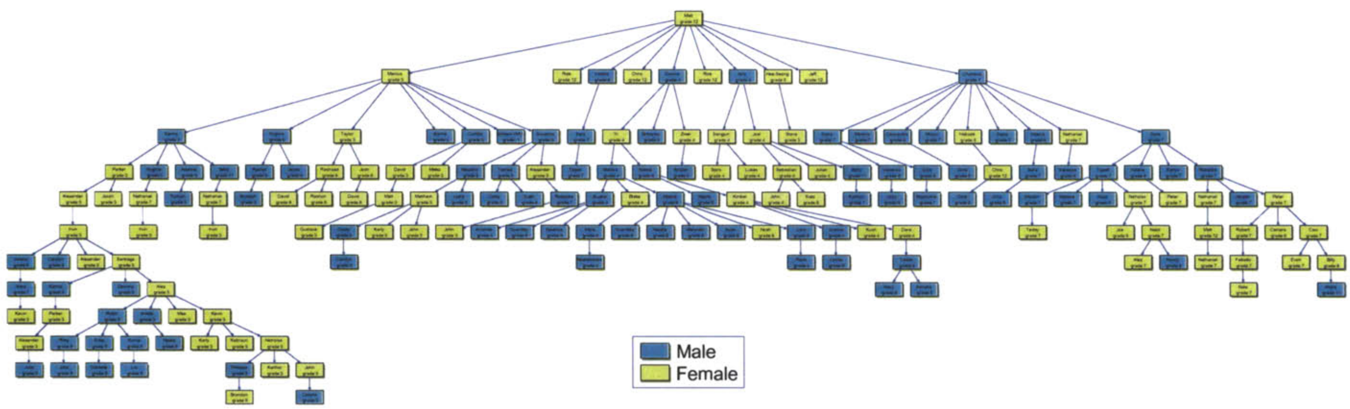

Figure 5-7 "Romance" I-Ball, Colorized by Gender

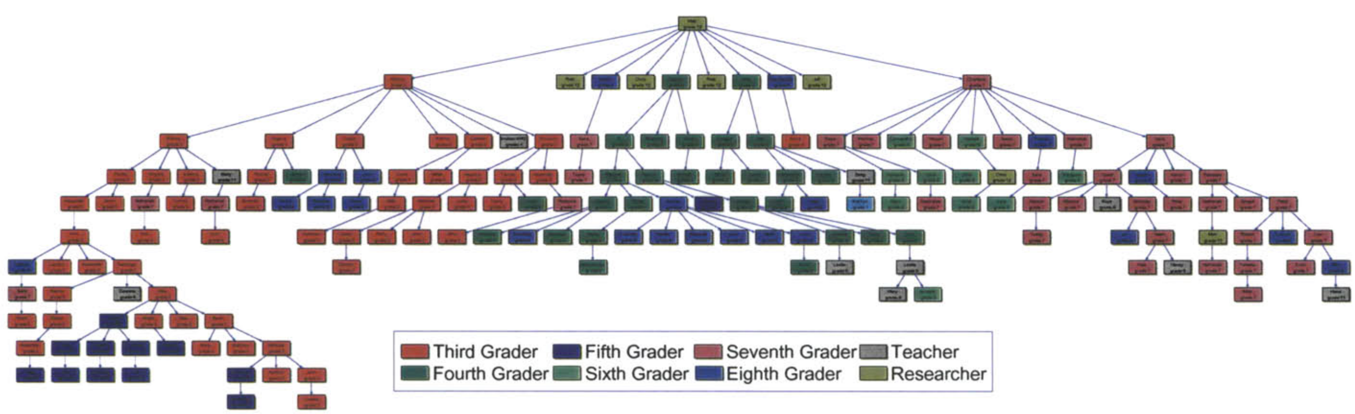

Figure 5-8 "Romance" I-Ball, Colorized by Grade

also saw this graph in Chapter 4). This graph shows that although the "Romance" i-ball is seen by roughly an equal number of males and females, it tended to get passed between people of the same gender, as evidenced by the clustering of blue and yellow nodes. The graph in Figure 5-8 shows the same type of insularity within school grades – i.e., people in a particular grade tended to pass it along to other people in that grade.

Collectively, colorized i-ball network graphs like the ones shown have the potential to broaden how people think about the complex nature of diversity. The graphs show that diversity should not just be measured in terms of the relative size of populations, but also in terms of their interaction. Also, diversity should not be conceived of solely in terms of single dimensions, such as ethnicity. Moving beyond these narrow definitions, the graphs show that while the school looks quite diverse in terms of race, it looks like a very traditional school in terms of segregation by grade. Hopefully, future users of Folk Computing technology will be able to use these types of graphs to build a better sense of their community in terms of its diversity.

5.2 New approaches to teaching traditional subjects

We believe that many subjects can be made more interesting and accessible to students when they are introduced in a more ecologically integrated fashion. The following sections present some lessons we learned about this, presented in the context of the variety of subject domains we explored.

5.2.1 Spelling

A Kindergarten teacher used our technology to help teach her class how to spell. She designed a quest-like game where kids used handheld devices that displayed pictures of spelling words, such as "cat" or "sun." In order to spell the word, the child needed to collect the appropriate letters in sequence from other kids who had those letters in their first names. For example, to spell "cat", a child would first have to connect their device to someone with a "c" in his or name, such as "Jack," then find another person with an "a" in their name, etc. This turned the normally sterile spelling drill into a much richer activity that was well integrated with the kindergarten ecology. The spelling quests were very popular with the kids, and the teacher made use of them for several weeks.

Figure 5-9 Kindergartners Playing a Spelling Quest

There was a reciprocally supportive relationship between the spelling i-ball activity and the kindergarten ecology: the i-ball activity preserved the ecology by not building too much structure on top of it. The ecology then played an important role in animating the i-ball activity. Figure 5-9 shows a boy engaged in a spelling quest collecting a letter he needed from a classmate. This photograph nicely highlights how this activity integrated with three key dimensions of the Kindergarten ecology.

Social dimension: The spelling quest made spelling an interpersonal activity. Finding a letter meant figuring out who you knew that had that letter in their first name. When it turned out the first person that came to mind was absent that day, the kids had to think a little harder, and maybe approach someone who they did not know as well. As is the case with all Folk Computing technology, the activity gets "read" in terms of the prevailing the social reality: people's affinities for and knowledge about each other add energy and meaning to the endeavor (Borovoy, Martin et al. 1998).

Temporal dimension: In Figure 5-9, the boy playing a spelling quest gets a letter from a girl who is engaged in another activity (drawing a picture). The i-ball activity did not require the simultaneous participation of everyone in the class. The i-socket's small size and wearability made it possible for most students to be pursing other activities, and still be available at short notice to provide a letter from their name to someone in need. Technology, especially technology designed for a group, is not always as willing to disappear into the background.

Spatial dimension: Children used the whole kindergarten space for the spelling quest activity, seeking out kids in the sandbox out back, and then going over to the names written above each cubbyhole to figure out from whom they could get their next letter. Moving around and between all the people and things in a kindergarten environment gave the spelling quests a nice physical quality.

5.2.2 Science

The seventh grade science teacher wanted to include the i-balls in a unit she was teaching on the periodic table. Together, we came up with a few activities that we thought might work. Both activities required that we associate a particular element to each person. For the first activity, students created element quests (e.g. "Find an element that is a noble gas") and gave them to each other to solve. The other activity, which seemed a little more challenging, was a guess-the-rule game: kids created "Dr. Science" i-balls and programmed them so that the i-ball would only jump to someone else if Dr. Science was "interested" in the type of element. When another student was given one of these i-balls, she had to try to figure out the category of elements Dr. Science was interested in (e.g. radioactive elements or heavy metals) based on who he would jump to.

Our attempt to integrate the i-balls into the science curriculum was not nearly as successful as the spelling activity. In retrospect, we realized that we failed to respect the key role that the teacher plays in supporting a classroom ecosystem. With the spelling activity, the teacher played the lead role in its development. Even though spelling was not on our list of what could be learned with i-balls, the teacher was able to use built-in functionality to make the spelling quest i-balls. The understanding she developed of the activity through designing and building it served her well when she had to deploy it in her class. To support the science i-ball activities, we had to make some modifications to the underlying i-ball software. Although the changes were small, they led to our being more centrally involved in the subsequent i-ball and activity design, and then in the deployment of the activity with the kids. Unfortunately, as the teacher later said to us, we didn't have the knowledge of her classroom's dynamics, or the relationships with the students that would have helped smooth the introduction of the activity. And she didn't have enough knowledge about the i-ball activity to play the role she normally did. The result was an awkward, overly structured activity that seemed to poison an otherwise vital classroom dynamic. The science teacher said, and we agreed, that while she loved the i-ball activity at the school, she did not like what happened in her class.

There was one more minor, but important flaw in the science activity. The assignment of a particular element to each student felt forced. We had initially thought this association was going to be made in the weeks before the i-ball trial, and that students would have done substantial research into these elements already. The schedule changed, however, and students wound up having a very weak identification with their elements. This served as a lesson on the importance of integrating technology into the "deep features" of an ecology (e.g. kids' names in the spelling activity) to insuring the activity's meaningfulness.

5.2.3 Data Analysis and Mathematics

Both sixth and seventh grade classes used the i-ball visualizations – particularly the ones involving bar, line, and pie charts – as an opportunity to practice their data analysis skills. In particular, the seventh grade teacher was interested in helping her students prepare for the MCAS assessment exams the following year, where they would be required to "identify and evaluate patterns and trends in data," "draw conclusions based on data or evidence presented in tables and graphs," and "based on research findings, identify alternative hypotheses or explanations of relationships among variables" (Education 1998). To this end, her students engaged in some interesting discussions of issues such as "Why were most of the ten most popular i-balls created by boys," where they used the graphing tools to test a variety of hypotheses including "Most of the participants in the trial were boys" and "Most of the i-balls were created by boys". The teacher, who was initially cynical about the relevance of the i-ball project to her educational goals, was very enthusiastic about the way these conversations went.

Although we did not get a chance to try this with students, we were able test for ourselves how i-balls might be useful for exploring some powerful ideas about randomness. On our own, we created two hitcher-type i-balls. The first one, called "shorter1", was programmed to automatically jump from its current i-socket to another it is connected to if the other's owner is shorter. In this way, it will seek out the shortest person it comes in contact with. We programmed the second hitcher, called "shorter2", to do the same thing, with the additional behavior that every so often it would jump no matter whether the other person was taller or shorter. The jumping rules from the two hitcher programs, and the resulting diagrams of the hitcher's travels through the school from the i-ball visualization tool, are shown in Figure 5-10.

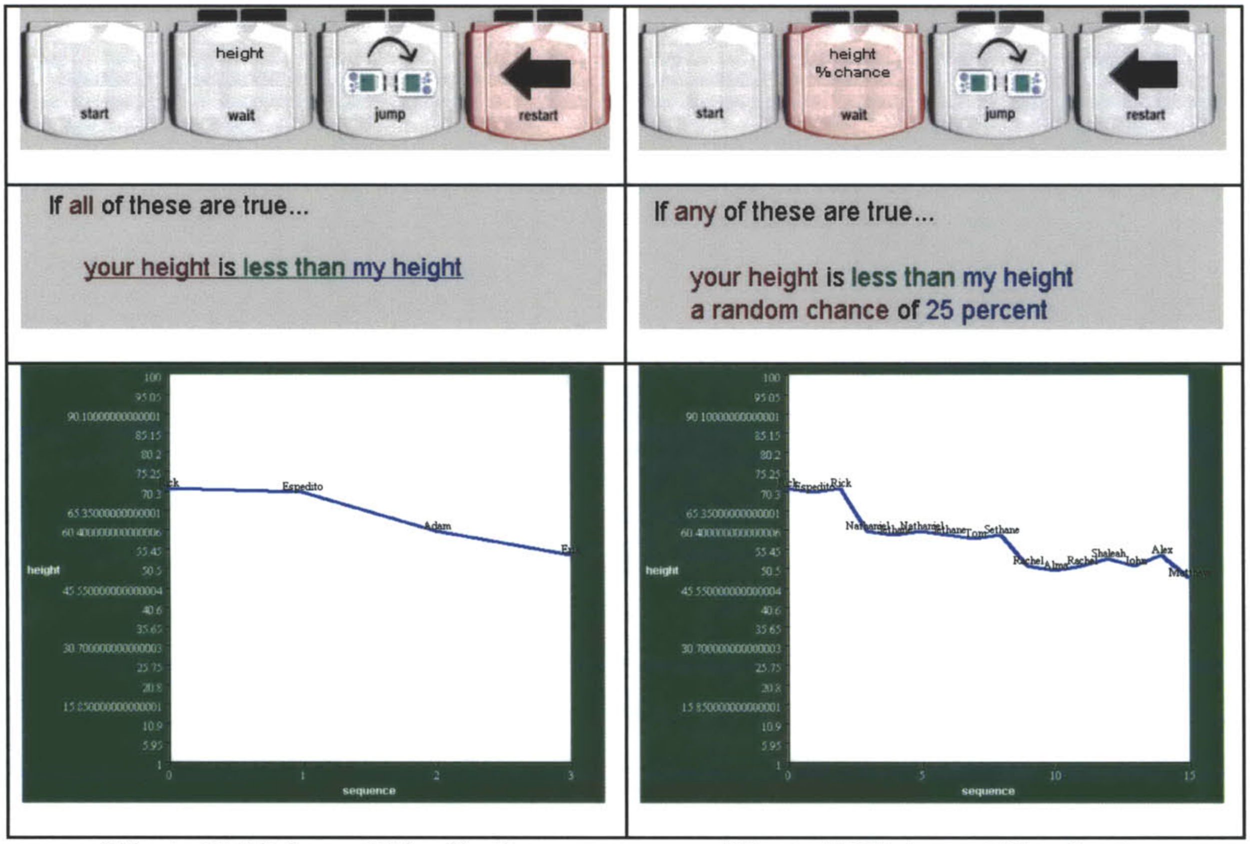

"Shorter1" Hitcher and Visualization

"Shorter2" Hitcher and Visualization

Figure 5-10 Variations of "Shorter" Hitcher and Resulting Graphs

Figure 5-10 Left shows the 4-block graphical hitcher program on top the top layer, the contents of the rule in the "wait" block in the middle layer, and the visualization of the hitcher's travels through the school in the bottom layer. The line graph in Figure 5-10 Left foes from left to right and shows the height of each person that the Shorter1 hitcher caught a ride on. As one would expect from the Shorter1 rule, the hitcher travels to shorter and shorter people monotonically. In contrast, the graph in Figure 5-10 Right shows that the Shorter2 hitcher, which has some randomness built into it, moves both up and down along the height dimension as it jumps from person to person, although the trend is downward. It turns out that Shorter2, the one with some random motion, ultimately winds up on a person 6 inches shorter than Shorter1, and only three people away from the shortest person in the trial. We think these types of experiments that the students could run have the potential to be excellent lessons in deep ideas involving probability, optimization, and the positive and counterintuitive role that noise can play in bouncing a system out of a local minimum. The fact that the subjects of this study will be friends and acquaintances of the experimenter, and that the experiment might involve variables as significant as height is to a child, suggests that such experiments would be engaging, as well.

5.2.4 Computational Ideas

In a manner similar to research on online learning communities (Bruckman 1998), i-ball technology created a "computational culture" where hundreds of students and teachers learned to write computer programs as a way to entertain themselves and their friends, to influence their community, and to explore their own ecology. Furthermore, there was some evidence that participants learned how to program from each other, either through direct interaction with each other or by making use of each other's i-balls as teaching examples. In future work, we would like to further explore the great potential that ecology of learners has for this type of peer-to-peer learning.

6 Conclusion and Future Work

Our exploration of Folk Computing has served to identify an important but largely ignored research problem: how to help people build a sense of community in informal, dynamic face-to-face settings. Folk Computing aims to solve this problem by facilitating the construction, circulation and visualization of resonant texts. This thesis has presented an evolution of Folk Computing technology and theory, as well as the results of several substantial user trials that establish the technology's promise for helping people explore and experience their community ties. At the conclusion of this stage of the Folk Computing research, there are two clear next steps. First, a more complete ethnographic study of Folk Computing technology should be undertaken to complement the formative research detailed in this thesis. Second, a much larger Folk Computing experiment should be performed to investigate the potential of this technology to build a sense of community on a more global scale. These projects are explored in the next two sections.

6.1 Designing an Ethnographic Study

On the first day of the school trial, we spent the morning inside the school's computer lab handing out i-sockets and giving people a brief introduction to technology. Although there was some excitement around the i-balls, we were disappointed that there was less "buzz" than at the first i-ball trial at the conference. At 1 PM, we took a lunch break and walked outside the building. Suddenly, we saw all these kids running around on the playground, eagerly showing and sharing their i-balls with each other. One ran up to me, thrust his i-socket toward my face and said, "Have you seen this i-ball? It's famous!" At this point, the activity had been running for only three hours.

Unfortunately, although the activity blossomed outside the more rigid structures of the computer lab, we as researchers were stuck inside keeping the whole system operating. This meant we were not able to observe how people used and made sense of the i-balls to the extent we would have liked. Although there were ways we could have done things differently, I also believe that we were dealing with an ineluctable paradox involving the study of Folk Computing prototype technology. This is important enough to explore briefly, even at the risk that it sounds like a rationalization.

As we have discussed previously, Folk Computing technology cannot be tested in small, controlled laboratory conditions; it is too hard to create the requisite amount of informal, unstructured, "background" type interaction. In fact, almost all the known failures of Folk Computing technology – when users failed to be energized by it – have occurred when people have attempted to use it in small, overly structured settings as a foreground activity. For example, someone tried to encourage a group of advertising executive to test out the Thinking Tags while sitting around a conference table.

Apparently, the CEO said something about it being silly, and then everyone else put theirs down. As discussed in Chapter 5, we did some unsuccessful experiments with a science class at the school where the i-balls were used as a foreground activity. Finally, consider the slow start of the i-balls that first day in the compute lab. It was not until the technology moved out into the wider school community that it began to take off. Like folklore, the technology is designed to occupy the interstices of the social structure.

The paradox is that prototype Folk Computing technology cannot be studied within a traditional laboratory, but it can't be thoroughly studied outside the lab either. The resources required to deploy Folk Computing prototypes in a community with enough complexity and dynamism to be worthy of study leaves insufficient resources to mount a careful ethnographic study of such a complex community. Even if there were enough resources, it would be very hard to do an appropriate study design without having collected some data and gained some experience from a prior trial of sufficient scale. In the end, one must conduct a large-scale formative experiment, which is what we have done with each of the Folk Computing technologies. Now we are in a position to design a more complete study. The following paragraphs briefly outline an ethnographic study of the i-ball technology.

A central claim of the i-ball technology is that it helps people build "a sense of community" – a phrase that means both a feeling of community or fellowship, and an understanding of community dynamics and structure. We will focus initially on the first meaning. Our principle means for gathering data on this in the formative trial was through interviews and observations conducted during and after the trial. These interviews were somewhat fruitful. For example, one of the teachers reported during a post-trial interview that the school community had been hurt by series of events in the previous two years, and that she felt the i-ball activity brought the community together in a very healing way.

Unfortunately, we were only able to gather these types of observations on a frustratingly random and sporadic basis. This erratic source of data contrasted with the near-complete computational record of i-ball creation and circulation. This put us in the role of paleontologist, with access to a very complete fossil record reflecting the who, what and where of people's activity, but only a small amount of data relating to why. For a more careful study, we would like a means of capturing more of the vast array of human-level thought, intention and interaction that corresponds to this dense computational record.

Two different teacher-initiated i-ball activities suggest the potential of content analysis of student work as a means for collecting a broader set of human-level data that would reflect on participants' sense of community (Robson 1993). The four images in Figure 6-1 were created by four different kindergarten students asked to write something about the i-ball activity. A cursory analysis of these images hints at some different conceptions of the project. For example, the person who wrote the upper-left document was thinking about the i-balls in terms of his own body when he said "todoy i got my ibol and it was col and they are not the cind of ibols that are in are hed " (sic). In contrast, the two people who did the lower-left and upper-right pictures as objects disconnected from themselves or anyone else. Finally, the document in the lower-right reflects the author's concept of i-balls as a social activity. It has two unique qualities: not only

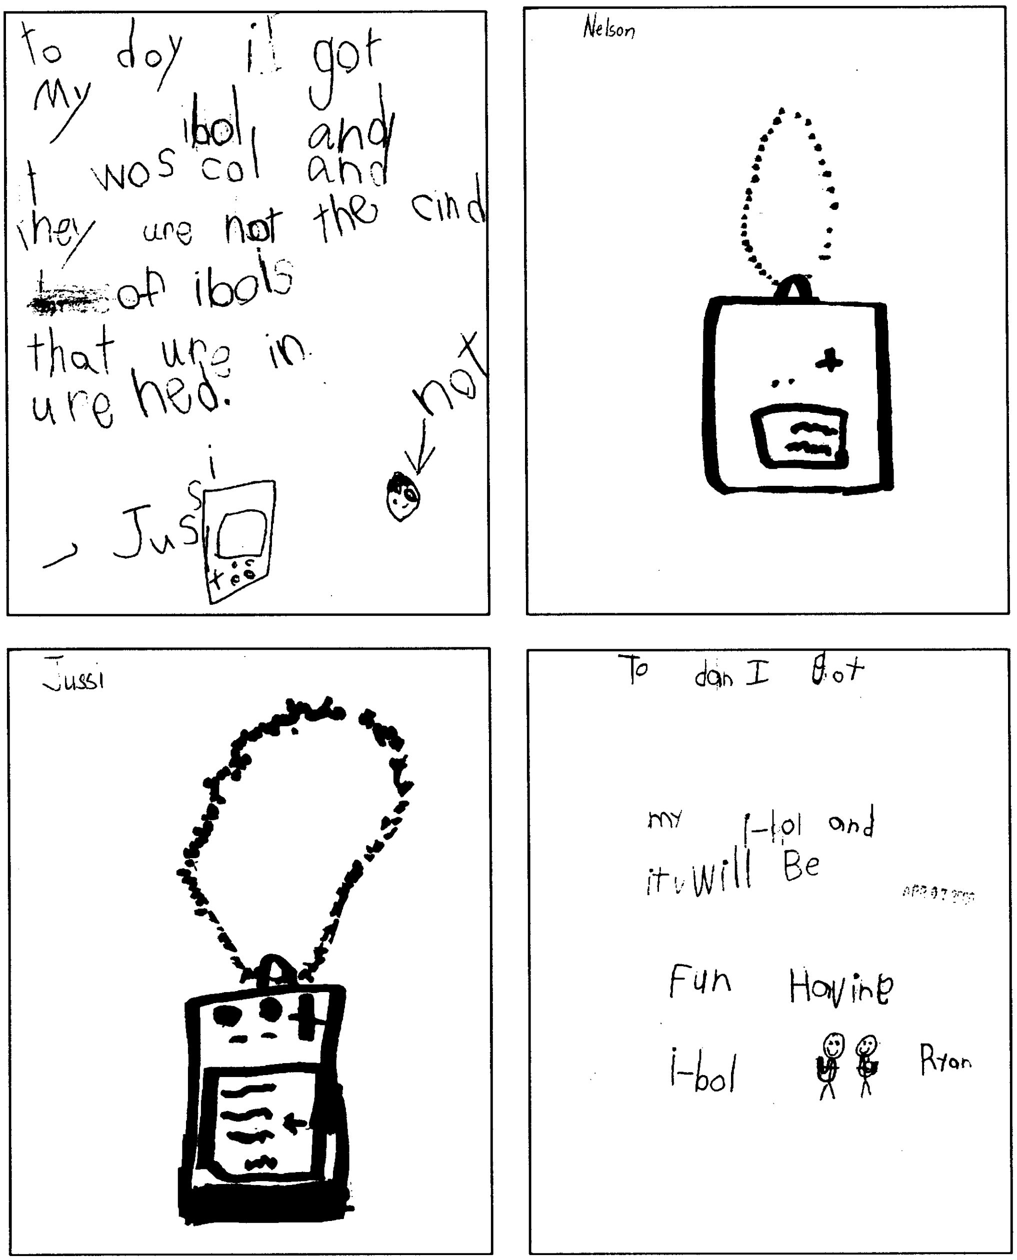

Figure 6-1 Kindergarten Drawings of the I-Ball Activity

is an i-socket shown being worn by a person, but there are actually two people shown interacting with their i-sockets and with each other.

The school librarian gave us another method for potentially collecting a larger amount of ethnographic data about how people were using the i-balls. She was in charge of the school's "television news" program. Every morning a group of students delivered a newscast comprised of live and taped segments to the school over closed circuit TV. The librarian asked one group of students to put together a segment on the i-ball activity. Their work on this illustrated the potential for equipping many kids with video cameras to document their i-ball experience in a future trial. This could provide a decentralized, unobtrusive means for gathering data on the large number of human-level interactions that went on in times and places where only kids were present.

Finally, a survey instrument could be useful for determining the role that the i-balls played in helping people develop a community feeling. We did a preliminary version of such a survey in the i-ball conference trial, where we asked people an open-ended question about their reaction to the i-ball activity. Some of the responses that related to community are quoted in Figure 6-2. A more carefully crafted survey could give more detail about the role the i-balls played in building community.

> "I found it very stimulating as ideas were exchanged between young and old… The i-balls also had a lot to do with this."
>
> "In trying to solve the quests I have never seen my son so willing to go up to people he did not know to seek information. By the end of Saturday he was engaging in conversation with anyone and everyone."
>
> "My favorite part was that it was a good way to get a sense of who was in MindFest also and also to get to now (sic) some more new people."
>
> "This also acted as an ice breaker between strangers. It was as if you had the right to ask them questions. … I believe the I-ball activity enhanced the MindFest by helping introverted people interact with extraverted people quicker."

Figure 6-2 Survey Results from I-Ball Conference Trial

One thing a survey will not detect is whether the community building effects of the i-ball activity are due to the novelty of the technology or to something deeper. Indeed, there is some reason to question this, since the data shows that i-ball activity substantially decreased during the second week of the school trial (see Figure 6-3). This could be due to the fact that two weeks was enough time to reduce the novelty effect of the technology, but not enough time to produce a strong "familiarity effect." In other words, the types of i-ball activities that we introduced became less exciting, but there was not enough time for most teachers and students to appropriate the i-ball technology for their own interests. A longer-term study would help explore whether the i-balls could outlive their novelty.

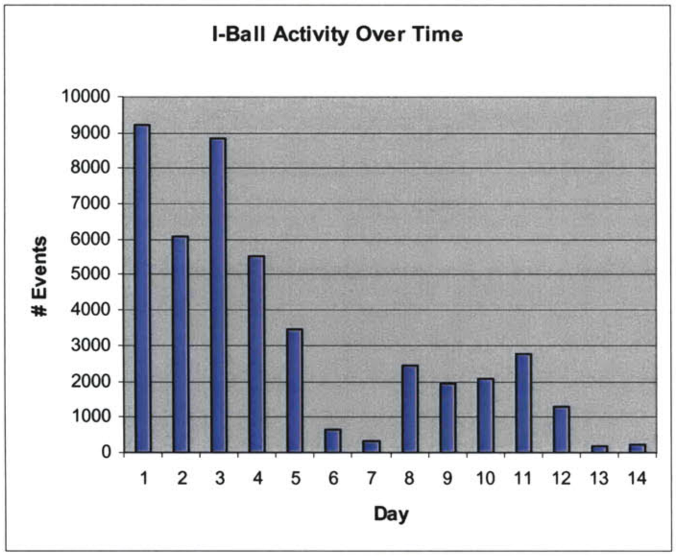

Figure 6-3 Total Number of I-Ball Events At School Over Time

What about a study designed to assess the i-balls' success in establishing the second type of "sense of community" – an understanding of community dynamics and structure? Chapter 5 provided several examples of students' insights about the structure of their community recorded while they were viewing different visualizations. To collect this type of data in a more systematic way, we could build a science curriculum unit around the type of social inquiry enabled by the i-balls. In this context, we could administer pre-tests and post-tests to see how students' concepts of community structure evolved over time.

Finally, in addition to assessing the change in users' sense of community, a more formal study should examine whether this change effected a change in behavior. For example, it would be interesting to know whether an increased awareness of insular behavior brought about by the types of colorized i-ball network visualizations in Figure 5-6 through Figure 5-8 caused people to attempt to interact with more people outside of their gender, grade or ethnicity. One could test this by staging the introduction of various types of visualizations during the trial, and then looking for changes in the subsequent interaction data.

6.2 Designing a Larger Folk Computing Trial

We believe that Folk Computing activities get more interesting the larger and more heterogeneous the target population. Therefore, we would like to push Folk Computing to the next scale. If it was interesting for third graders to see how their i-balls had spread into higher grades, imagine how exciting it would be to see that an i-ball one had made had spread from person to person all the way to Japan. The following sections sketch three different scenarios for a scaled-up Folk Computing activity.

6.2.1 Phosphori

This activity would give kids a sense of community on a global scale.

Before each Olympics, a small group of runners carry a torch from Athens to the host country. We would like to add a collection of digital torches, carried by millions of kids all over the world. The kids would carry the digital flames on their wrists, in specialized Olympic watches called "Phosphori" (ancient Greek for "torch bearers" or "light bringers"). To pass the flame, kids would simply bring their wrists together, and the flame would jump from one watch to another. The ultimate goal: to pass the flames all the way to Athens for the 2004 Olympics.

There would be five different Torch Watches, each one resembling a different colored ring on the Olympic flag. To create a new flame, a group of five kids wearing the five different color watches must put their hands together (see Figure 6-4 Top-left). A new flame glows on the watch faces, along with a unique URL (see Figure 6-4 Top-right). On the Web, kids could add pictures and messages to "their" flames -- and see visualizations of all of the flames (see Figure 6-4 Bottom-left). Each flame, over time, would represent a community of kids from around the world.

There are games associated with the flames. Kids gain points each time they pass a flame; with enough points, they can start a new flame. Flames must be cared for; if they are not passed frequently, they die out.

During the opening ceremony, the lights will dim, and hundreds of children in the stadium will hold up their hands, their watches glowing red. These children (chosen for the key roles they played in the torch relay) will each be carrying one of the flames that made it all the way to Athens. They will then run down onto the field, and one by one, will touch their watches to a large digital torch display (see Figure 6-4 Bottom-right). An enormous animated flame will ignite on the screen, where each piece of the flame will be an image of a child who carried it, like an animated version of the "Photomosaics" created by Rob Silvers (see www.photomosaic.com).

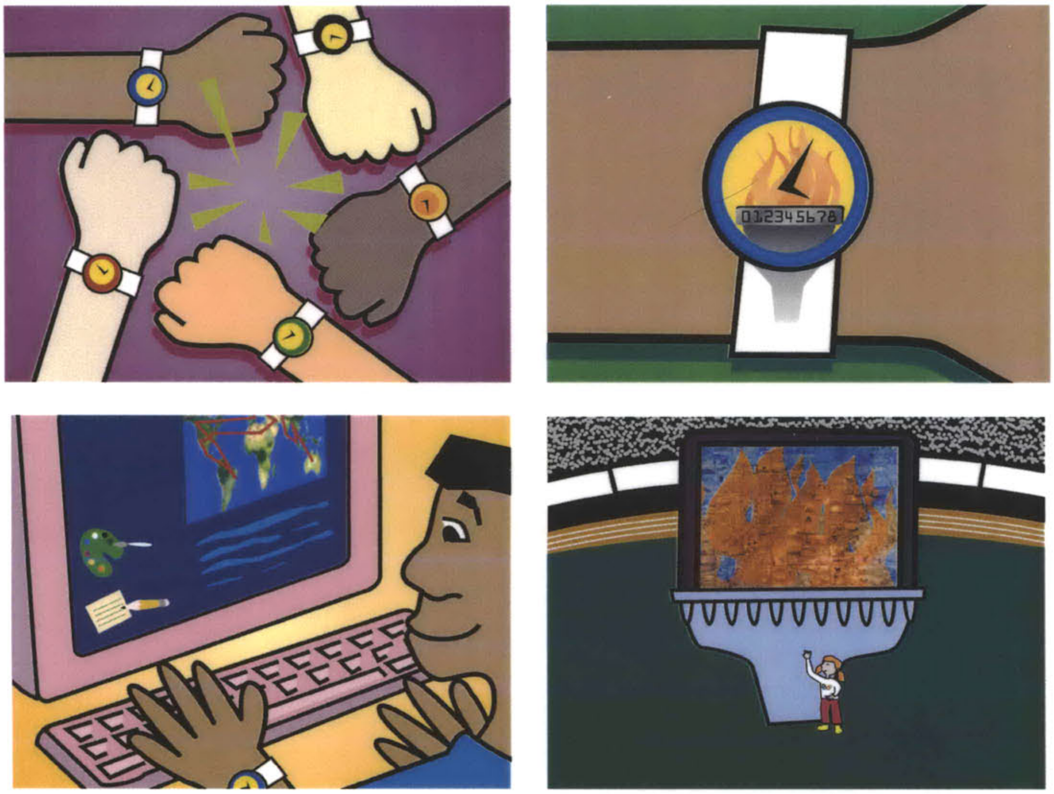

Figure 6-4 The Birth of a Flame (Top-left), A Unique Burning Flame (Top-right), Tracking and Annotating a Flame (Bottom-left), and Opening Ceremonies Flame Visualization (Bottom-right)

Drawings by Michelle Hlubinka

6.2.2 MemeMail

Although the focus of Folk Computing has been on face-to-face community building, we are also interested in whether visualizing the circulation of online folklore – in the form of frequently forwarded emails – might also help people build a sense of community. To begin exploring this, we have built an early prototype, called MemeMail, which allows people to create augmented email messages that provide anonymous, demographic information about who else has read them. Like a kind of visual "cc:" field, this information is meant to help people get a sense of the community circumscribed by a particular message. This might help create the kind of mutual knowledge around frequently-forwarded email that is usually missing from this medium, since – unlike an article in the New York Times – one can not easily get a sense of how widely or to whom a piece of email lore has circulated. MemeMail does this while still preserving the decentralized, democratic nature of frequently forwarded email as a medium.

Figure 6-5 Top shows an example MemeMail that actually circulated in Spring, 2001. It calls for a "voluntary rolling blackout" to protest President Bush's energy policies. When people open this message up in their email inbox, in addition to seeing the message text, they see the blue map of the United States that to provides an approximation of a "live" view of where people who have read this email live.[14] By looking at this visualization, the knowledgeable email reader can in fact get an interesting sense of the community within which it is circulating. Comparing the MemeMail map to the map of the 2000 Presidential election results (from http://www.sweetliberty.org/images/ mandatemap.jpg) shown in Figure 6-5 (which was not included in the email) hints that this message is mainly circulating in Democratic strongholds. In other words, it is preaching to the converted.

One of the major virtues of MemeMail as a research platform is the low cost to participation, compared to one of the hardware-based Folk Computing technologies. With the MemeMail prototype, it should be possible to do a controlled study to test the effect of Community Mirror visualizations on people's behavior. We could send out a particular MemeMail to a random group of people. As a control, we could then send out a plain email with the same text as the MemeMail, but without the Community Mirror, to a separate random population. We could then see whether the

[14] Producing a reliable and accurate visualization like this is challenging. These messages are html-enhanced, and all email readers handle html slightly differently. Also, in this implementation, the MemeMail server figures out the reader's location by analyzing their IP address, which is not that accurate. A more accurate but more invasive method would be to ask for people's zip codes the first time they receive a MemeMail.

Email subject: Fwd: Roll Your Own Blackout

ROLL YOUR OWN BLACK OUT
THE FIRST DAY OF SUMMER
JUNE 21, 2001 THURS EVE,
7-10pm, all US time zones

In protest of George W. Bush's energy policies and lack of emphasis on efficiency, conservation and alternative fuels, there will be a voluntary rolling blackout on the first day of summer, June 21 at 7pm - 10pm in any time zone in the USA (this will roll it across the country).

It's a simple protest and a symbolic act. Turn out your lights from 7pm-10pm on June 21. Unplug whatever you can unplug in your house. Light a candle to the sungod, kiss and tell, make love, tell ghost stories, do something instead of watching television, have fun in the dark.

Forward this email as widely as possible. Use the map below to see what areas of the US haven't gotten this message yet, and send it to people you know there. Write to your government representatives and environmental contacts. Let them know we want global education, participation and funding in conservation, efficiency and alternative fuel efforts -- and an end to over exploitation and misuse of the earth's resources.

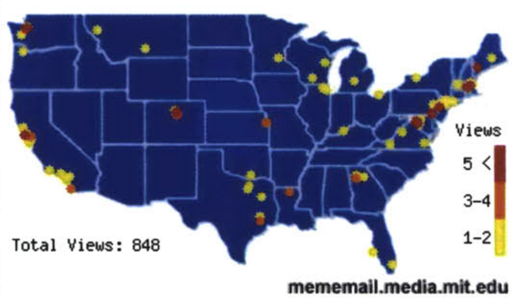

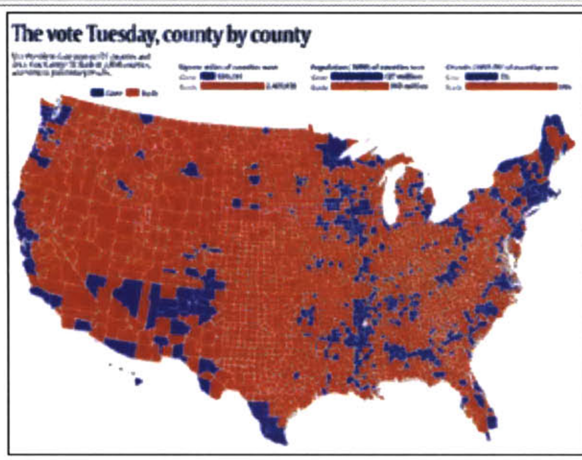

Figure 6-5 A Political MemeMail Example (Top) and County-by-County Results of the 2000 Presidential Election (Bottom)

presence of the visualization effected how the email was forwarded. For example, with the above MemeMail, people might start to see that the message is mainly circulating in Democratic regions, and make more of an effort to forward it to their Republican friends.

Bibliography

Anonymous (1998). Audience response systems: Where to get them. Presentations.

Borovoy, R., F. Martin, et al. (1998). GroupWear: Nametags that tell about relationships. CHI '98 Conference on Human Factors in Compuing Systems, Los Angeles, CA, ACM.

Bronner, S. J. (1986). Folk Ojbects. Folk Groups and Folk Genres. E. Oring. Logan, Utah, Utah State University Press.

Brown, J. S. and P. Duguid (1996). "The Social LIfe of Information." First Monday **1**(1).

Bruckman, A. (1998). "Community Support for Constructionist Learning." Computer Supported Cooperative Work **7**: 47-86.

Bruckman, A. and M. Resnick (1995). "The MediaMOO Project: Constructionism and professional community." Convergence **1**(1): 94-109.

Brunvand, J. H. (1981). The Vanishing Hitchhiker. New York, London, W. W. Norton.

Cashion, J. (1986). "'You'll love this one...': An exploration into joke-prefacing devices." Journal of Language and Social Psychology **5**(4): 303-312.

Cassell, J., J. Smith, et al. (1999). CrossTalk. Siggraph 99 Catalog on Electronic Art and Information, Los Angeles, CA, ACM.

Chesley, H. (1991). Asynchronous background netowrking on the Macintosh. Develop.

Clark, H. H. (1996). Using Language. Cambridge, Cambridge University Press.

Clark, H. H. and T. B. Carlson (1982). Speech acts and hearers' beliefs. Mutual Knowledge. N. V. Smith. London, Academic Press: 1-36.

Clark, H. M. (1992). Arenas of Language Use. Chicago, University of Chicago Press.

Colella, V., R. Borovoy, et al. (1998). Participator simulations: Using computational objects to learn about dynamic systems. Summary Proceedings of CHI '98, Los Angeles, ACM Press.

Danesh, A., K. M. Inkpen, et al. (2001). Geney: Designing a Collaborative Activity for the Palm Handheld Computer. Proceedings of CHI, Conference on Human Factors in Computing Systems, Seattle.

Davis, F. (1992). Fashion, Culture and Identity. Chicago, University of Chicago Press.

Dawkins, R. (1989). The Selfish Gene (2nd Edition). Oxford, Oxford University Press.

Dean, K. (2000). Al, Dubya Caught in Web of Satire. Wired Online.

Denison, D. C. (2001). Words to Live By: Off-Base Net Dialogue Now Belongs to Many. Boston Globe. Boston: A1.

Donath, J. S. (1995). Visual Who: Animating the affinities and activities of an electronic community. ACM Multimedia 95, San Francisco, CA, ACM.

Douglas, M. (1999). Implicit Meanings: Selected Essays in Anthropology. London and New York, Routledge.

Dourish, P. and S. Bly (1992). Portholes: Supporting Awareness in a Distributed Work Group. Human Factors in Computing Systems, Monterey, CA, ACM.

Dundes, A. (1989). Folklore Matters, The University of Tenessee Press.

Education, M. D. o. (1998). Guide to the Massachusetts Comprehensive Assessment System: Science and Technology. Malden, MA, Massachusetts Department of Education.

Emerson, J. P. (1969). "Negotiating the serious import of humor." Sociometry **32**(2): 169-181.

Erickson, T. (1996). "The world wide web as social hypertext." Communications of the ACM **39**(1).

Fish, S. (1980). Is There a Text in this Class: The Authority of Interpretive Communities. Cambridge, MA and London, Harvard University Press.

Goffman, E. (1959). The Presentation of Self in Everyday Life. New York and London, Doubleday.

Goffman, E. (1963). Behavior in Public Places: Notes on the Social Organization of Gatherings. New York, The Free Press.

Goodwin, M. H. (1990). He-Said-She-Said: Talk as Social Organization among Black Children. Bloomington and Indianapolis, Indiana University Press.

Hancock, C. M., J. J. Kaput, et al. (1992). "Authentic inquiry with data: Critical barriers to classroom implementation." Educational Psychologist **27**(3).

Hanson, R. (1990). Could gambling save science: Encouraging an honest consensus. Eighth International Conference on Risk and Gambling, London.

Hofferth, S. L. and J. F. Sandberg (2001). Changes in American Children's Time, 1981-1997. Children at the Millennium: Where Have We Come From, Where Are We Going? T. J. Owens and S. L. Hofferth. New York, Elsevier Science. **6**.

Hollan, J. and S. Stornetta (1992). Beyond being there. Proceedings of the ACM CHI '92 Conference on Computer-Human Interaction, Monterey, CA.

Jason, H. (1997). Enhancing Your Presentations wtih the OptionFinder Audience Polling System. Ogden, UT, Option Technologies, Inc.

Jenkins, H. (1992). Textual Poachers: Television Fans and Participatory Culture. London, Routledge.

Kane, T. R., J. Suls, et al. (1977). Humour as a Tool of Social Interaction. It's a Funny Thing, Humour. A. Chapman and H. Foot. Oxford, Pergamon.

Knapp, H. and M. Knapp (1976). One Potato, Two Potato: The Secret Education of American Children. New York, W.W. Norton.

Koschmann, T., Ed. (1996). CSCL: Theory and Practice. Mahwah, New Jersey, Lawrence Erlbaum Associates.

Krauss, R. and S. Fussell (1990). Mutual knowledge and communicative effectiveness. Intellectual teamwork: Social and technological foundations for cooperative work. J. Galegher, R. E. Kraut and C. Egido. Mahwah, NJ, Lawrence Erlbaum Associates: 111-145.

Lynch, A. (1996). Thought Contagion: How Belief Spreads Through Society. New York, Basic Books.

Martineau, W. H. (1972). A Model of the Social Functions of Humor. The Psychology of Humor. J. H. Goldstein and P. E. McGhee. New York and London, Academic Press.

McCabe, N. (1998). "Glory, glory hallelujah, teacher hit me with a ruler: Gender and violence in Subversive Children's Songs." Studies in Popular Culture **20**(3).

McCrone, J. (2000). You buzzing at me? New Scientist. **165:** 20.

McDowell, J. H. (1979). Children's Riddling. Bloomington & London, Indiana University Press.

Mechling, J. (1986). Chapter 3: Children's Folklore. Folk Groups and Folk Genres: An Introduction. E. Oring. Logan, Utah, Utah State University Press: 91-120.

Mieszkowski, K. and K. Quistgaard (2000). Cell hell. Salon.

Milgram, S. (1967). "The Small-World Problem." Psychology Today **1**(May): 60-67.

Nunamaker, J., A. Dennis, et al. (1991). "Electronic Meeting Systems to Support Group Work." Communications of the ACM **7**(July): 40-61.

Opie, I. and P. Opie (1959). The Lore and Language of School Children. Oxford, Clarendon Press.

Oring, E. (1986). Chapter 1: On the Concepts of Folklore. Folk Groups and Folk Genres: An Introduction. E. Oring. Logan, Utah, Utah State University Press: 1-22.

Oring, E. (1986). Chapter 2: Ethnic Groups and Ethnic Folklore. Folk Groups and Folk Genres: An Introduction. E. Oring. Logan, Utah, Utah State University Press: 23-44.

Papert, S. (1980). Mindstorms: Children, Computers, and Powerful Ideas. New York, Basic Books.

Resnick, M. (1999). Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds. Cambridge, MA, MIT Press.

Resnick, M. (2001). "Thinking like a tree (and other forms of ecological thinking)." To be published in the International Journal of Computers for Mathematical Learning.

Robson, C. (1993). Real World Research: A Resource for Social Scientists and Practitioner-Researchers. Oxford, UK and Cambridge, MA, Blackwell.

Rogoff, B. (1994). "Devleoping understanding of the idea of communities of learners." Mind, Culture, and Activity **1**(4): 209-229.

Sack, W. (2000). Discourse Diagrams: Interface Design for Very Large Scale Conversations. Hawaii International Conference on System Sciences, Persistent Conversations Track, Maui, HI.

Shannon, C. (1948). "A mathematical theory of communication." Bell System Technical Journal **27**: 379-423 and 623-656.

Shardanand, U. and P. Maes (1995). Social information filtering: Algorithms for automating "word of mouth". Human Factors in Computing Systems, Association for Computing Machinery.

Soloway, E., C. Norris, et al. (2001). "Log on education: Handheld devices are ready-at-hand." Communications of the ACM **44**(6): 15-20.

Starner, T., S. Mann, et al. (1996). "Augmented reality through wearable computing." Presence, Teleoperators and Virtual Environment **6**(4): 386-398.

Stefik, M., G. Foster, et al. (1988). Beyond the Chalkboard: Computer Support and Problem Solving in Meetings. Computer-Supported Cooperative Work: A Book of Readings. I. Greif. San Mateo, CA, Morgan Kaufman Publishers,**:** 335-366.

Sussman, N. M. and H. M. Rosenfeld (1982). "Influence of Culture, Language, and Sex on Conversational Distance." Journal of Personality and Social Psychology **42**(1): 66-74.

Sweller, J. (1994). "Cognitive load theory, learning difficulty, and instructional design." Learning and Instruction **4**: 295-312.

Toelken, B. (1979). The Dynamics of Folklore. Boston, Houghton Mifflin Company.

Velde, V. d. (1997). Co-Habited Mixed Reality. Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, Aichi, Japan.

Viegas, F. and J. Donath (1999). Chat Circles. CHI 99 Conference on Human Factors in Computing Systems, Pittsburgh, PA, ACM.

Viluoja, E. (1996). Manifistations of the revenant in Estonian folk traditiion, Institute for the Estonian Language, Estonian Folklore Archives.

Want, R., A. Hopper, et al. (1992). "The Active Badge Location System." ACM Transactions on Information Systems **10**(1): 91-102.

Wasserman, S. and K. Faust (1994). Social Network Analysis: Methods and Applications. Cambridge, Cambridge University Press.

Webster (2001). Merriam-Webster Online. **2001**.

White, T. and D. Small (1998). An interactive poetic garden. Proceedings of the Conference on CHI 98 Summary: Human Factors in Computing Systems, Los Angeles, CA, ACM Press.