Visualizing Information Spaces to Enhance Social Interaction

by

Rebecca Wen Fei Xiong

B.S. Electrical Engineering and Computer Sciences, University of California, Berkeley, 1994
S.M. Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1996

Submitted to
the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2000

Author...
Department of Electrical Engineering and Computer Science
May 15, 2000

Certified by ....
Randall Davis
Professor of Computer Science and Engineering
Thesis supervisor

Certified by ....
Judith S. Donath
Assistant Professor of Media Arts and Sciences
Thesis supervisor

Certified by ....
Albert R. Meyer
Hitachi America Professor of Engineering
Thesis supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Student

Visualizing Information Spaces to Enhance Social Interaction

by
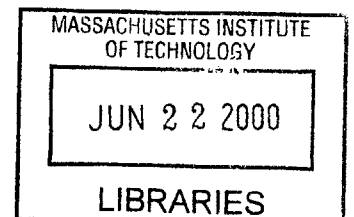
Rebecca Wen Fei Xiong

Submitted to
the Department of Electrical Engineering and Computer Science
on May 15, 2000, in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

Human beings are social animals. In real life, we are constantly in contact with others, whose activities help us decide where to go, what to see, and who to talk to. In information spaces like the World Wide Web (WWW), however, people cannot easily gain a sense of being with others, determine patterns of activities, nor find appropriate others to interact with. Our solution, LiveWeb, visualizes the underlying information structure and overlays the dynamic real-time presence of people. This live visualization serves as a basis for more effective interaction. LiveWeb has been deployed on a public Web site with over ten thousand hits a day. Experiments with real users of the system over several months have shown a strong preference for the presence display.

Thesis supervisors:

Randall Davis, Professor of Computer Science and Engineering
Judith S. Donath, Assistant Professor of Media Arts and Sciences
Albert R. Meyer, Hitachi America Professor of Engineering

Acknowledgement

I am greatly in debt to the many wonderful people who have made this thesis possible. I would like to thank Professor Judith Donath for her guidance and support for much of my thesis work. Her original and insightful comments added greatly to this thesis. She has also opened my eyes to ways in which effective designs can not only please the eye but also enhance our social interaction.

I would also like to thank my other co-advisors. Professor Albert Meyer has shown long and continuous support for my work, from its infancy, through various changes and many advisors, to its current incarnation. He has generously listened to my hopes and frustrations and offered invaluable advice and encouragement. Professor Randy Davis has lent his clear thinking and sharp insight to many of the problems I faced in the last months of this thesis.

I have been a member of two research groups: Computer Graphics Group at the Lab for Computer Science and the Sociable Media Group at the Media Lab. At the Graphics Group, Professor Teller has given me the freedom to pursue my own research and the machine and office space to pursue it in. Eric Brittain has helped me design and code the beginning of LiveWeb and has always lent an ear when I needed someone to bounce ideas off of. At the Sociable Media Group, Roy Rodenstein has offered me many thoughtful comments on my work (and equally thoughtful discussions about movies, a shared hobby of ours). I would also like to thank many other members of both groups.

I would like to thank Steven Drucker, Marc Smith, and others from the Virtual Worlds Group at Microsoft Research for making NetViz possible. Dean Alberta Lipson helped me formulate questions for my surveys. Elizabeth Harris from the Writing Center not only corrected the grammar and wording of this thesis, but also acted as my "superego" to help me think through the writing.

My friends have made my years at MIT not only challenging but also enriching. I'd like to thank Kathy Liu (my ex-roommate), Rama Mukkamala (my dinner buddy), Eric Wu (my baking buddy), Chi Nguyen, Brett Bochner (for the coffee hour debates), Elaine Shi, Gloria Lin, Bei Wang, Frank Wang, Victor Lee, Rosanna Tse (Go Bears!), Tom Chou (for introducing me to the enneagram, which made possible the experiments on LiveWeb), Edmund Yeh, Ruilin Zhao (for his great stock tips that I never took), Vernon and Beth Ingrams (for making Ashdown a welcome home for so many of us), Damian Isla, Janey Hoe (who helped me decide to come to MIT), Xiaowei Yang, Marina Bers (my jogging buddy), Mark Foltz, Will Neveitt, Derek Chiou (for his great taste in food), James C. Hoe (for his jokes), Sandeep Chatterjee, Dan Engels, and Ed Wang.

My mother has always encouraged me to pursue my dreams and has never wavered from her belief that I will fulfill them. For this and so much more I will always be grateful. My brother Wen has also been supportive through the years.

Finally, I would like to thank R.R. for everything he does and is.

Mission

This thesis was conceived during the writing of another thesis. While working on my master's degree about an algorithm for fast rendering of large 3D models, I spent hours navigating through city models on the computer screen. Although the city streets looked pretty and ultimately I could navigate at a fairly decent speed, I quickly found the task to be rather dry. What I missed, I discovered, was the ever-changing stream of people in a real-life city.

This led me to think about the lack of presence of other users on the Web. The Web, at that time, was largely for information browsing. Even today, while chat-rooms abound, we have not addressed the problem. I feel we may be in danger of building a different kind of pretty-looking model that lacks the richness of real-life city traffic. This thesis is the result of my search. Its mission is to demonstrate the following:

- Users prefer being able to see this presence of other users;
- Such presence can be visualized in a meaningful manner; and
- Such visualization can enhance users' Web experience.

To my mother, whose courage and strength will always be an inspiration to me.

Contents

1. Introduction

"What attracts people most, in sum, is other people, " concluded William Whyte of his pioneering study of urban street life [Whyte88]. He found that public areas are essential to a city's public life: they enable people to mingle with crowds, talk to friends, and watch strangers. Jane Jacobs, author of one of the most influential books on urban planning, also observed that streets and neighborhood parks allow people of diverse backgrounds to meet and socialize at a public level [Jacobs61]. Architects have pointed out that, "Being human, human beings are what interest us most. In architecture as in literature and the drama, it is the richness of human variation that gives vitality and color to the human setting" [Raskin60].

Today, a different kind of public street life is emerging. On-line environments, such as the World Wide Web, have become increasingly popular. At the end of the year 1998, the US alone had 60 million people accessing the Web, or more than one in five [Headcount98]. This has created new opportunities for interaction with users from around the world. Web-based chat rooms, message boards, and other internet tools like ICQ, AOL Instant Messages, and buddy lists have proliferated: in on-line environments as well as in real life, people long to reach other users. In fact, studies have shown that, in virtual worlds, people want the connection to other people more than any other experience [Heeter92, Lombard97].

Imagine the possibilities: a reader can see what the rest of the world, or maybe college students like himself, is reading from one moment to the next. A PC buyer configuring his purchase on-line can learn about the options offered by communicating with others making similar selections. During the holidays, an on-line shopper may see more people in the shopping sites and even run into friends or family members. Web-surfing can potentially be an enriching social experience. By allowing users to interact easily, the Web can become a more enjoyable environment for remote education, work, or community building.

Currently, however, the Web is still far from reaching its potential as an environment for interaction. There are Web-based message boards and chat rooms for interacting with other users, but they are not well integrated with the main

activity of browsing for information. Thus, even with millions of users on the Web, Web-surfing is still often a solitary experience.

1.1 Deficiencies of current on-line environments

Let us explore this problem in more depth. Comparisons to the real world can help us identify the deficiencies of the virtual world.

Lack of social presence

Real-life environments have a wonderful ebb and flow of activity. During the course of a day, a popular store may get a few visitors in the morning, then get very busy with shoppers during the day, and gradually quiet down as the evening comes. This ever-changing flow marks the passage of time and enlivens the otherwise rather static city scene [Lynch72]. Such lively flow is missing from the Web today. Users can interact with each other at designated areas, but the vast other parts of the Web, however information-rich and visually appealing, is essentially a ghost town.

This lack of presence means that Web users cannot easily perform a basic human behavior, reacting to others [Aronson54]. They cannot observe user patterns and behaviors that might give suggestions about how they should act. The Web today lacks the sense of a crowd gathering at a mall because of a sale, or of a new restaurant attracting more business than its neighbors. Web surfers cannot easily gauge if one Web page generates more user interest than another, and cannot act appropriately in response.

Lack of social identity

Traditional face-to-face interaction is often conducted with a limited number of people, so one will usually get to know his interaction partner in some depth. With virtual environments like the Web, one has a much larger number of people with whom he may interact. These environments do not have physical constraints of distance or traditional links such as employment in the same company. The people one encounters on-line, therefore, could have very different background from people he knows in real life. In addition, because on-line environments offer more choices of groups to interact with, he may have less time to spend in each group.

14

What the user needs are new ways to construct his own social identity and comprehend those of others. Instead of facing bland and faceless crowds, he needs to be able to determine potential partners for interaction. The various on-line environments currently available should help the user to quickly assess the different mix of people within them.

Lack of spontaneous interaction

Currently, the Web lacks means for users to connect easily with others. Users must go to a chat room to find each other. The analogy would be making the effort to go to a bar to meet and talk to others. Most interactions in real life, however, happen spontaneously as a part of everyday life, rather than being pre-planned. One can run into colleagues in the hallway, people-watch in a café, ask a passer-by for direction, or join a crowd in cheering. On-line environments lack the "casual, public contact at a local level -- most of it fortuitous" [Jacobs61] that make up a healthy public life.

1.2 Importance of addressing these deficiencies

In some sense, our effort is similar to city planners' search for ways to enliven a city. In traditional village environments, neighbors often have similar backgrounds, know each other well, and form friendships based on proximity. With the Industrial Revolution and the creation of cities, a large number of people from diverse backgrounds began to live in close quarters. While cities have built impressive, imposing skyscrapers as symbols of their technical skills and wealth, their residents have increasingly faced problems of alienation [Wirth38]. Gradually, city planners have come to emphasize the importance of a healthy public street life. They have explored possible solutions, such as by designing city plazas and common areas that aim to encourage people to gather and interact with each other [Whyte88].

In a sense, virtual environments like the Web extend the growth of cities by connecting people from around the world. While we make larger servers, faster networks, and more powerful search engines, we also need to create a more social Web that connects rather than distances people [Oldenburg89, Rheingold93]. Otherwise we may make the same mistakes big cities in the past have made, building glittering but cold, impersonal virtual environments.

1.3 Our approach: visualization for enhancing interaction

How can we overcome the barriers to effective on-line interaction? Trying to reproduce the real world is obviously infeasible, and perhaps the effort is misguided. The virtual world has many virtues of its own. Here, people are connected, whether they are conscious of it or not, by informational objects, such as the Web pages they are viewing. This allows for new kinds of crowds, patterns and movements. Unlike the movements in the real world, user movements in virtual environments can be captured explicitly. Using the captured data, we can show users the presence of other users within the environments, and make explicit the shared interests among them.

1.3.1 Visualizing information structure

To display presence within an information space effectively, we must first consider how to show the structure of this space itself. Real environments have spatial structure and containment. For example, a person can be inside an office, which is inside a building, which is on a street. Users are familiar with many cues in their physical environment. Offices have doors that can be left open or closed to show different levels of accessibility for interaction. On the road, the traffic can be heavy or light, indicating different degrees of activities. What about virtual environments? What structure should we use to show presence in a Web site?

Visualizing information structure is a research topic in its own right. Though our work makes use of techniques from that area of research, our focus is on helping users to understand the flow of people and their interaction within the information space. This means that we emphasize

- First giving broad overview, then supplying more details on demand
- Showing relationships between different informational objects, rather than just the objects themselves.
- Exposing patterns and anomalies

The NetViz visualization to be discussed in Chapter 5 illustrates how to show relationships among different topics or different discussion groups. Showing such relationships can help users to identify topics or groups of interest. WebFan, presented in Chapter 4, structures messages from different discussions to illuminate user traversals within the discussion site.

1.3.2 Visualizing social context

In addition to visualizing information structure, our approach visualizes the social context within the virtual environments. By social context we mean both the sheer presence of users and their patterns of activities. By showing presence, we help people to gain a sense of being with others, and create a stronger sense of social connection. By showing people's activities, we allow them to better distinguish their fellow participants and identify people they might want to interact with. We show

- Real-time presence

- Individuals' interaction history

- Relationships among people

- Relationships among people and objects

WebFan shows presence of users on a message board through the course of a day. PeopleGarden, to be discussed in Chapter 6, also visualizes a message board, but focuses instead on its participants and their posting history. Because some posters make large amount of uninteresting postings, the amount of response from other people at the board is a more revealing measure of the level of interaction with other users than posting history alone. NetViz shows clearly just who talks to whom in a discussion. It also indicates who participates in a particular discussion, exposing more about users' interests. We will see more examples later of using these visualizations to determine behavior patterns and identify people of interest.

1.3.3 Interaction ramp-up

Our definition of interaction incorporates both traditional two-way interaction between two or more people, and one-way interaction where one person may, by specific actions or more passive presence, provide information that others make use of. Users of on-line environments need an easy means for expressing spontaneous opinions or gradually increasing their level of interaction. At the present time, users on the Web have to go intentionally to chat rooms to interact with others. Certain bulletin boards or newsgroups actually discourage people from posting short "me too" messages showing support for ideas. These practices create a barrier to interaction. Meanwhile, the Web continues

to offer more topics and areas of interaction for a user, each vying for his attention. He has therefore less time and effort for possible interaction.

What users need is the ability to interact at different levels. For some Web sites, a user may simply want to identify "hot spots" of people, the Web pages that are very popular. For other sites, he may be curious to find out more about other participants. In other cases, he may actively look for people to interact with, hoping to sustain interaction over time. This means that users may need a whole continuum of interaction support, from general trend- or people-watching, to gaining interest in interaction, to actually interacting. Our approach is two-fold. First, we propose better integration of Web browsing and interaction, so that the barrier to interaction is lowered. Next, we propose the interaction ramp-up, outlined in Chapter 3, to offer a range of interaction support.

1.4 Bringing presence to the Web

Our goal is to bring the liveliness of real life to information spaces, but tailor it to the unique needs and affordances of the on-line world. To create a practical and comprehensive solution requires assembling all the parts discussed above. What is required in a system that provides social presence and allows for more informed interaction? What should it look like? How should the user interact with it? We have constructed a system that incorporates a visual map of the site, the display of users' presence, and chatting capability. Challenges arose not only in thinking through the functionality and interface, but in deploying the system on public Web sites and tailoring to the needs of their audience. The LiveWeb system will be discussed in detail in Chapter 7.

Our work lies in between traditional engineering theses and more recent design theses. Traditional engineering theses describe new ways of solving a known problem faster and better. Design theses offer insights and thoughtful designs that explore what is possible. Our work lies in between, in that we aim to identify and articulate what we feel is an important new problem, and propose a practical solution that works with existing technical capabilities and human requirements. What does this combination imply for the evaluation of our work? For engineering theses, if the problem it addresses has not been solved, presenting a solution is sufficient; otherwise, one can compare the performance of the new solution to existing solutions. Design theses can be evaluated on the originality of insights

and the depth of thinking around the problem. We will evaluate our work by taking an intermediate approach: we will demonstrate that the problem we have identified is indeed important and that our solution is sufficient.

The best way to test our claims is through actual use. By building the LiveWeb system, we have provided a solution to the problem of lack of presence on the Web. How actual users interact with the system and show preference for its presence display can demonstrate both that the problem is real and that the solution is satisfactory. Chapter 8 details the experiments we have conducted using the LiveWeb system. Our experimental results showed that users indeed preferred the presence display, by at least three to one. Our survey results suggest that users were able to understand the display and actively used it to gain information about other visitors to the site.

1.5 Areas of previous work

Our work is influenced by research from diverse areas: classic works in social psychology and urban studies, as well as recent works in information visualization, computer supported collaborative work, and collaborative virtual environments. Insights from **social psychology** guide our thinking about how to create effective self-presentation and interaction. **Urban studies** provide hints about how to create environments conducive to social interaction. Our visualization is a form of **information visualization**. However, our emphasis is not on visualizing large amounts of data for traditional data analysis, but on presenting patterns and behaviors about users to help them to interact. **Computer supported collaborative work** has investigated user awareness and lightweight interaction within an organization. We are interested in extending this work to the Web, where there are a larger number of users with diverse backgrounds and interests. Finally, **collaborative virtual environments** also deal with on-line crowds and interest management, but these simulate three-dimensional real spaces rather than more abstract information spaces. These are also typically set up to pursue a common task or explore a space rather than for interaction. We will describe these related researches in greater detail when discussing different aspects of our work.

1.6 Thesis outline

This thesis contains three parts, reflecting the hybrid nature of this work.

I. Theory – A survey of previous work in social presence and our framework for interaction ramp-up

II. Design – NetViz, WebFan, and PeopleGarden, and their design rationale

III. System and Experiment – The LiveWeb system, integrating several designs and chat, and evaluations conducted using the system

Part I, Theory, discusses the background and theoretical framework for our work. We have pinpointed the lack of presence as an important problem with the Web today. Chapter 2 defines social presence and surveys existing work on the topic. Chapter 3 lays out our thinking on interaction ramp-up, our model for user interaction.

Part II, Designs, presents several of our experiments for creating effective visual designs, and discusses their design goals, rationales, and the lessons learned. Chapter 4 describes WebFan, which visualizes presence on a simple Web-based message board. Chapter 5 presents PeopleGarden, which explores how to meaningfully represent users for on-line forums. Chapter 6 discusses NetViz, which visualizes the relationships among people and informational objects.

Part III, System and Experiment, puts our thinking to practice. Chapter 7 presents LiveWeb, a real-time presence system that we have developed and deployed on a public Web site. LiveWeb visualizes the underlying information structure of the site and the social presence within it as a basis for interaction. Users are able to see the presence of other users, identify potential interaction partners, and interact with them. Chapter 8 describes the experiment we have conducted using LiveWeb and the quantitative and qualitative feedback we have received from actual users of the system. Finally, Chapter 9 reflects on the evaluation results we have obtained and their implications on our theory, and summarizes lessons learned from our designs.

Part I. Theory

The theory section discusses ideas that have guided our thinking and have led to our designs. We proposed in our introduction that the lack of social presence is a barrier to interaction on-line. Chapter 2 starts with a discussion of social presence, addressing questions such as

- What is social presence?

- How can it affect an information space?

- What are ways by which social presence can be supported?

We will examine previous works from several related fields. What makes our work different is our focus on the presence of users in an information space, rather than presence itself. The presence of other people affects how we see the information; where people go in the information space in turn affects how we view them.

The intimate relationship between information and people is really the basis for our work. Most chat rooms provide little context for interaction. We propose using the informational content as the context for interaction. Chapter 3 addresses interaction based on informational space:

- How does interaction come about in real life?

- How do interactions in real life and on-line differ?

- What is a workable model for interaction in an information space?

2. Social Presence

We define social presence and discuss its impact on information space. Social presence in information space has many benefits: it enables social navigation, allows people to connect to others, and enriches their interaction with others. We discuss support for social presence in information space by examining previous solutions in the real and the virtual worlds. We will examine more closely this support along the axes of time, degree of user involvement, and levels of information about users.

2.1 Definitions

How can we make the Web a more social environment? To better address this problem, we need to understand the connection between social interaction and information space. We start our exploration by defining both of these components.

2.1.1 Information space

By information space we mean a collection of information objects. Real-world examples include a set of lecture notes or a warehouse of books. We focus here on virtual collections: examples include a file directory containing multiple files, a message board with messages from various people, a Web site with multiple pages, or even the entire World Wide Web. Note that the connections among the different information objects can vary. The messages on a message board are bound in a message-reply hierarchy, whereas a Web site may contain many links among its pages, and sometimes has no clear hierarchy. There may also be other, less well-defined relationships among objects. For example, organizing messages by topics can result in overlapping subsets. Thus, there may be multiple ways of presenting the information space to a user, depending on his or her interest. A hierarchical view can help user determine more clearly what he has explored, while a topic-centered view may be more useful for browsing.

As more and more companies put up Web sites and more people create homepages, we will have a greater repository of on-line information. We will also spend more time in these informational spaces as part of our daily lives. The challenge, for both the designers of these information spaces and the society at large, becomes, How can

we make these information spaces more interesting? How can we make them fulfill not just our thirst for more information but also our need for connecting with other human beings?

2.1.2 Social presence

Social presence is a less well-defined concept. Traditional communication theorists have applied the term as a measure in comparing different communication media, particularly telephone and email. The Social Presence Theory proposes that the amount of presence depends on the interaction of a medium and task at hand. Presence is measured with a set of bipolar scales such as impersonal-personal, unsociable-sociable, cold-warm [Short76]. The Media-Richness Theory measures a medium's social presence based on its capacity for immediate feedback, number of senses, personalization, and language variety [Rice92]. The absence of social cues in a particular medium can affect, positively or negatively, how we comprehend, assess, and interact with other users. For example, people who are less physically attractive or less quick to react found emails to be liberating: on-line interaction allows them to be judged more on the basis of their ideas and thereby gain a larger voice in their organization [Sproull95].

The recent development of virtual reality has created a new vision of the future and changed the focus of discussion on social presence. Carrie Heeter distinguishes it from personal presence, defined as the amount and depth of sensory stimuli a medium provides, and from environmental presence, which measures how much the environment responds to a user. To her, social presence is the effectiveness of the medium in creating the illusion of being in the medium with other people. How much other people acknowledge and interact with a user determines how real the experience feels to the user [Heeter92]. Frank Biocca writes

> *The minimum level of social presence occurs when users feel that a form, behavior, or sensory experience indicates the presence of another intelligence. The amount of social presence is the degree to which a user feels access to the intelligence, intentions, and sensory impressions of another.*

Heeter's and Biocca's definitions relate closely to our work. We are not interested in the sheer amount of sensory input or perceived realism a medium provides, but focus instead on its capacity to facilitate substantive interaction with a large number of others. A medium can provides clues about the underlying social environment, such information about the number, location, and actions of other users. By illuminating overall trends and individual behaviors, a medium can enable users to interact more effectively with each other.

24

2.1.3 Social Presence in Information Space

Now that we have defined social presence and information space, let us see how these two can be integrated. First, we note that social presence can already be found on-line, but it is separated from the information content offered on-line. Existing chat rooms allow people to interact with others. However, most of these spaces are purely social: each room may have a theme or topic, but otherwise users are left on their own to find material for discussion. Thus, the conversation in many chat rooms consists of mostly greetings as people explore different chat rooms with different decorations and topics. In graphical chat rooms, people may comment quite a bit on the visual interface itself. The point is, there is little content to sustain conversation between people who do not know each other very well. Exceptions are controlled environments such as on-line games, where people have a clear goal. Having a human or robot host whose function is to create conversations can also enhance interaction.

Information spaces without presence are also possible. Digital libraries, such as the ACM digital library, contain large collections of articles from different sources [ACM99]. Although there might be thousands of people using the system at the same time, each of them is unaware of the others. The Web as a whole contains a large number of pages but very little social presence. (People's homepages may be their social representation on-line [Chandler99], but our research addresses mainly the possibilities for direct influence and interaction in the real-time presence of others.)

Libraries in real life, on the other hand, can exemplify information spaces with social presence. The collection of books forms the information space, and users create the social presence. During weekdays or weekends we may see different kinds of users. We can see which sections of the library are crowded and who goes to which sections. We can ask librarians for assistance with the information space, or we may run into people we know. Many libraries also serve a more explicit social function for the community of readers they have gathered. They provide information about local community activities that bring together people with similar interests. Various reading groups or children's activities also foster interactions.

There are also on-line examples of information spaces that reflect non-real-time social presence. Web-based message boards or Usenet newsgroups allow people to post their comments and respond to others. The set of all the posts forms the information space. Users can interact with others by responding to others' posts. However, people cannot communicate with each other in real time. They also have no way of knowing whether their message has been read by a few or many: those users who do not post are simply invisible to other users.

We address these issues by adding real-time social presence to these message boards and to the Web as a whole. We can then enable users to see others, understand their behavior patterns, and interact directly with them. First, let us understand better what such an endeavor could achieve.

2.2 Why add social presence to information space

Why should we augment information spaces, such as the Web, with social presence? Why would this integration be better than the current separation of Web browsing and chatting, for example? Think of a simple real life example of a trip to the museum. Going to a museum with friends is fun because we can share right away our thoughts about the paintings and statues we see. The feedback is immediate. It is also fun to note what other people find interesting about the exhibits, or hear a little of the official tour. Their reactions and comments may bring us new insights and enhance our museum experience.

Then think about going to an empty museum with no one around, looking at all the paintings by oneself, and trying to collect all one's thoughts together to discuss it with friends later. While this can still be fun, it lacks the liveliness of sharing with other people in real time. The same differences exist on-line, between an information space with social presence and one without. Let us discuss in more detail what benefits social presence on-line could bring.

2.2.1 Seeing where the hot spots are

One benefit of our approach is that the presence of others helps one identify information of general interest. Social navigation is the act of determining where one should go in an environment by observing the flow, concentration, and behaviors of others. Stanley Milgram found that there might be a human instinct to follow a crowd. He

conducted studies where a group of students would deliberately look up at a building, and found that over 70% of passers-by would follow the collective gaze [Milgram77].

Real life has many examples of social navigation: we may choose one restaurant over another because we observe more people there. The implication is that others may have information about which restaurant is better, information which we do not have. Similarly, if we are not sure how to get out of a subway station, one way is to follow others. These examples can illustrate the nature of social navigation. Following others often works, but not always. A group of people may be under the same misconception about the quality of two restaurants or have different preferences from us. A subway station may have multiple exits that lead to different places. We note two lessons here: 1) with a lack of information or in a new situation, it can be useful to follow other people; 2) the odds can be improved when more information is available about the people we see. For example, knowing that most of the people in the restaurant are there for the first time might change one's opinion.

The social navigation experience is currently missing from the Web. One cannot follow others to find Web pages of general interest nor avoid the well-trodden paths and look for new territories.

2.2.2 Connecting with other people

As the Web gains popularity, more people will spend more time on-line for information and entertainment. We can turn their solitary Web-surfing experience to a more social one by making them aware of the presence of others on the Web. Knowing that others are around can create a greater sense of togetherness, allow people-watching, and lead to more spontaneous interaction. The need for greater connection with others has long been studied. The sociologist Ray Oldenburg argued more than a decade ago for the importance of social connectivity, especially in today's suburban life, where people have become isolated and detached [Oldenburg88]. As Oldenburg points out, "community is not for everybody, but for many, it's necessary."

We see the same need for social connectivity in computer-mediated experiences. Carrie Heeter found in her study of virtual reality games that most people preferred to play them with other humans only (58%) or with mixed human/computer players (40%). Only 2% preferred playing with computer opponents alone [Heeter92]. She

27

concluded that people "want that connection [with other people] more than any other experience." Her finding is not unique. Two research labs were linked through video, with the intention of conducting impromptu research meetings remotely. Soon, however, people started to use the link informally to check who might be around [Bly93]. This result will not come as a surprise if one examines the history of technology adoption. The telephone started out as a device for business, and only later became adapted for social interaction [Sproull91]. Other technologies that can potentially connect people are likely to follow similar trajectories.

How can adding social presence to information space enhance people's social connectivity? In his studies of urban spaces, William Whyte found that people tend to socialize in the paths of busy activity, rather than away from it [Whyte88]. People often gather around doors and as a result block traffic. In these busy areas, there are more choices for what to do, and higher likelihood of running into acquaintances. Similarly, spontaneous on-line interaction is best conducted not in separate chat rooms, but alongside content pages that receive high traffic. By allowing users to become aware of each other and their respective activities, we can turn the Web into a rich social environment. People can watch what others do, react to others, and interact with each other.

2.2.3 Using information as a context for interaction

Judith Donath argued for a Sociable Information Space where the information can serve as "a context for community" [Donath96]. She noted that televisions already transmit cultural material to people, material that serves as a common ground for conversation. Unlike television programs, which bring people who meet through real-life connections closer, on-line spaces can promote interactions between people who may never meet in real life. Howard Rheingold, in one of the first books about the on-line life, gave a more personal account. He initially turned to newsgroups to gather information, but through his interaction with others, he gradually built up relationships. What ultimately kept him in the newsgroups was the connections he made with others in these groups [Rheingold93].

We can expand the effect of interaction beyond designated message boards or discussion forums onto the Web itself. By enhancing the Web with presence of others, we can leverage its vast informational content to bring a greater number of people together around different topics of interest. Oldenburg identified coffee shops, community centers,

and general stores as the "great good places". They serve as hangouts where people can mix with others from the same neighborhood, and as "sorting places" where one meet people with similar interests [Oldenburg88]. The Web, enhanced with presence of others, can become a new kind of virtual sorting place.

In describing the factors contributing to successful sites, Philip Greenspun has also pointed out two benefits: the giving and seeking of information and the ability for people from diverse physical locations to connect. The combination of these two factors can create what he calls a community: "a group of people with varying degrees of expertise in which the experts attempt to help the novices improve their skills" [Greenspun99].

Content that exists on the Web at present already attracts people and can serve as a context for discussion. In the course of discussion, people meet each other and establish relationships. They also generate new artifacts from their interaction, in the form of postings, traffic patterns, or chat messages that can augment the existing information. This greater repository of information, enhanced with social presence, can attract additional users to the space. The loop expands.

2.3 How to support social presence in information space

Having discussed the benefits of augmenting information spaces with social presence, we now examine how to achieve and sustain such enriched spaces. There are three sub problems:

1. Laying out content legibly

2. Showing the presence of other users

3. Encouraging user interaction

To address these problems, we first examine a range of solutions drawn from the real and virtual worlds. There are several related areas of work: urban planners and architects from the real world, and computer supported collaborative work (CSCW) and collaborative virtual environments (CVEs) from the on-line world. Real world solutions may contain the accumulated wisdom of many years. However, the on-line world presents new challenges: presence in the virtual world has fewer cues, and must be explicitly recorded, analyzed, and played back. Let us look at how some of the problems have been addressed by the different disciplines in these two worlds.

1. Laying out content legibly

Urban planners are interested in creating cities that people enjoy living in. They have explored ways of laying out a city so that it will be easy to navigate. In his seminal work *The Image of the City*, Kevin Lynch lists five elements that can help people navigate a city: nodes, paths, edges, districts, and landmarks [Lynch60]. We can apply these features in the design of virtual spaces that are easy to understand and navigate: the user can draw upon his experience of navigating in real environments [Chalmers96, Foltz98]. For example, people can find their way in virtual space by making reference to a given landmark, or by following a particular path between nodes. The main question for the designer is, what are the legible features of a given virtual space.

2. Showing presence

Awareness research in the area of Computer Supported Collaborative Work (CSCW) explores ways to better support a small group of people (typically less than 50) working together on a common goal, whether a short-term design project long-term collaborative research [Bly93, Ackerman95, Pedersen97]. CSCW researchers have found that the awareness of other people and their tasks plays an important part in the satisfaction and efficiency of ones' own work. Work in this area is similar to ours in the goal of supporting presence in a shared space. The difference is that we are dealing with much larger groups of people and less well-defined goals.

Research in collaborative virtual environments (CVEs), like our work, deals with the embodying of users in an on-line environment. CVEs typically are based on a 3D spatial metaphor. They could be used by a large number of people, and a crowd representation scheme might be needed. Our work, by contrast, needs not be constrained by a spatial metaphor and thus can have more freedom in representations.

3. Encouraging interaction

Urban planners have examined the issue of laying out cities for more effective use. Chris Alexander pointed out that cities should not have separate sections for distinct functions, but that the different functions should be mixed, so

that each neighborhood is a vibrant community. This notion ties in with our approach of integrating rather than separating browsing for information and interacting with fellow surfers.

Urban planners have also investigated how to revitalize a city by fostering interaction among its residents. For example, William Whyte has examined how to build plazas that will effectively provide outlets for human recreation in large cities [Lynch60]. He found that simple techniques, such as more seating space or easier access from the sidewalk, all contribute to a better-utilized plaza. These findings can help us think about how to build more accessible on-line environments, such as by making on-going chat visible to surfers everywhere on the Web site.

Architects have long wrestled with the problem of designing buildings that best meet the needs of inhabitants and withstand the passage of time. Buildings do not just provide a passive space; they can affect what people do and how people do it. The high ceilings and dark interior and stained glass windows of a church can inspire humility and reverence, while a spacious lobby, shared kitchens, and common areas in a dorm may promote interaction and socialization. Various affordances of an on-line environment can also have such strong impact. Requiring users to create an account to post a message, for example, may increase accountability, but may also discourage some new users with relevant comments from participating.

2.4 Showing presence over time

We have discussed some of the different perspectives on our problem that related disciplines could provide. We next examine in more detail, using examples from these disciplines, how the sense of presence can be supported through time. Past presence is the history of user activities. Showing where others have gone in the past can help current users decide where to go. Current presence is the representation of current users of the system. Knowing where they are can be the first step towards interacting with them. Future presence considers how to create an information space that can encourage more user activities. In his book *What Time Is This Place*, Kevin Lynch argued eloquently for the need to add time to the design of space [Lynch72]. Past, present, and future are intimately related values for people. Knowing a place has been well traveled in the past not only encourages current and future use but can affect how we value a place. The same thing can be said about an on-line environment.

31

2.4.1 Past: showing history of user activities

In real life, books wear and tear as they are read more often. A book may open easily to certain places that are read often. Its margins may contain written notes and parts of the text may be highlighted. This sense of history is lacking in digital objects. Will Hill et al. proposed the History-Enriched Digital Object [Hill93]. Their EditWear and ReadWear are simple scrollbars overlaid with histograms that show the reading history of various files [Hill92]. These scrollbars can be used to identify passages that people dwell on and to compare the reading and modification histories of multiple users.

Footprints is an example of history-enriched Web site [Wexelblat99]. People's traversals through a Web site are used to create a map of the site: the nodes represent individual Web pages, and the edges indicate popular paths through the site. Footprints can also suggest to users on a particular page what the popular next steps are. Some of these ideas are also found in PatinaMap, in which a geographic map becomes the interface to a Web travelogue [Schutte98]. Nodes in brighter shades represent regions that are more traversed; wider edges connecting the nodes represent greater amounts of traffic between nodes.

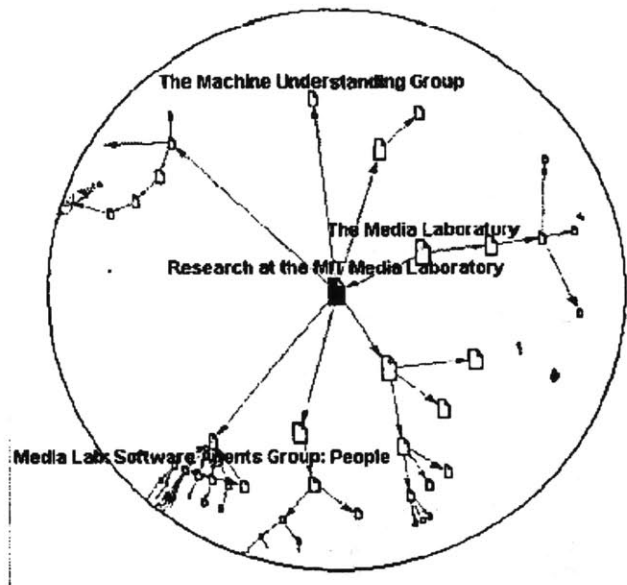Figure 2.2 A Footprints Map showing the connection between URLs traversed [Wexelblat99].

In summary, past presence can be represented by showing

- Where activities have occurred
- How much
- Which paths

The first two items help new users locate popular places that are likely to attract more people in the future; the third can help users find trails that are more likely to get them where they want to go.

2.4.2 Present: representing current activities

There are various ways of conveying a sense of presence in a space. Elin Pedersen suggests a framework that examines the different display schemes along the axis of abstract versus concrete [Pedersen97]. Bly et al. applied a concrete approach using multi-media video and sound for directly providing social awareness in shared workspace [Bly93]. Mark Weiser and others at Xerox PARC created a more subtle abstract display using dangling strings and windowpanes [Weiser91, Weiser96]. A less aesthetically pleasing but more precise representation is a simple histograms of how many people are currently using various parts of the system [Ackerman95]. This scheme was developed to attract new users for groupware systems, since the popularity of these systems depend on the number of users.

In attracting more users, it is less important to show the actual individuals. For chat rooms and virtual world, however, it is essential to show the actual participants. As illustrated in Figure 2.3, avatars can be used to graphically represent users. On can also augment an avatar with information about the user: his name, what he is saying, etc.. Benford et al. have explored using avatars to represent not just individuals but crowds in virtual
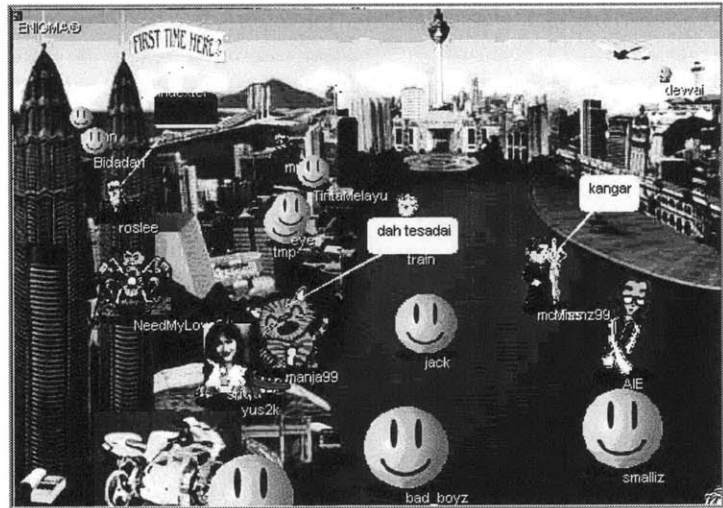
Figure 2.3 Avatars in graphical chat rooms show who is around and what they are talking about. Http://www.onlive.com.

environment [Benford97]. The avatar scheme has also been extended to showing presence on the Web. Judith Donath's 1010 Web site at the Media Lab allows users to see the recent visitors to the page, represented by pictures people select for themselves [Donath96].

These works show current presence by indicating

- Sense of presence of others
- Overall number of users
- Individual representation of current state

The sense of presence helps a user to become aware of sharing the space with others, while the actual numbers can attract new users to a system and encourage use. Avatars can further show individuals who are in a space and may be available for interaction. To find the right interaction partners, however, may require knowing more about others' interests, which is difficult to determine in existing systems.

2.4.3 Future: building for interaction

How can we build an information space that supports continued use and attracts future use? Lynch notes that "An environment that's precise and ordered may inhibit new patterns of activity" [Lynch60]. Stewart Brand, writing about livable buildings, offered similar suggestions that we might adopt: the low road and the high road [Brand94]. Adopting the high road means for the designer (who is often also the user) to thoughtfully create an environment that fits the user's purpose well right from the start. More relevant for us in creating an environment used by many others is the low road. Taking the low road means creating a low-budget, hack-it-yourself environment that allows the current users to adapt it easily to their needs. The low road provides the most basic functionality and flexible tailoring tools for users. The underlying notion is the acceptance that buildings will not remain as they are, but will change through the years, as shown in Figure 2.4.

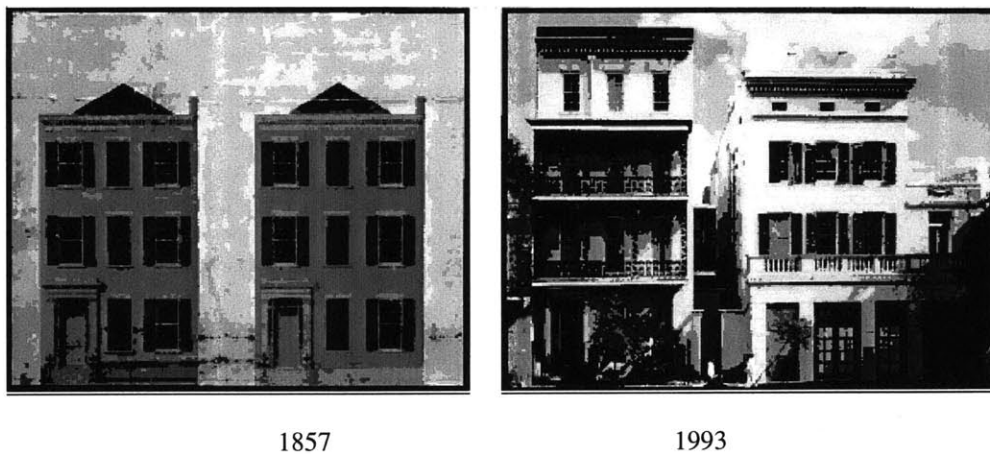1857   1993

Figure 2.4 The same two buildings more than a century apart [Brand94].

Let us examine how we can makes use of the low road approach in the design of virtual environments. One example of an on-line environment that shares the same "open" characteristic is Dourish and Belotti's ShrEdit program

[Dourish92]. It contains two types of editor windows: shared and private. In a shared window, what the user types will be seen by others, unlike in the private window. Users came up with different schemes of using the program. Some groups worked out a scheme such that each user maintains one shared window to show finished edits. Others share all the windows. Users can easily tailor the program to suit their purpose.

Designing for future use is by far the more tricky aspect of support over time. To ensure satisfactory future use, a system needs to allow user the ability to tailor the interface. William Whyte observed an example from real life: people in public places such as parks tend to move their chair just a few inches, even though the place they move it to may be no better. They need to feel a sense of control over the environment. The ChatCircles project provides a virtual example. Users can move their circle representations freely on the screen. In addition, the flexible scheme allows user to come up with surprising new ways of using the system. They can move their circles to form new groups (as originally designed), or play with convention of personal space by putting their circle close to others (not originally designed) [Viegas99].

We see that future presence can be encouraged by

- Providing the basics
- Giving more options for arranging the space

A conceptually simpler system will allow a user to manipulate the space more confidently. Giving users options for arranging their own space will then allow them to fit it to their needs.

2.5 Providing passive vs. active presence

There are two ways of obtaining presence. Real world objects sustain wear and tear in the natural process of being used. An old book may open at places where people have read it often, and reveal its history. A Web site can also maintain a record of use in the form of a Web access log. We call this type of data *passive presence* because users leave this record without consciously exerting effort.

In real life, readers can also actively highlight passages of a book or write down thoughts in the margin for themselves or others. This corresponds to leaving annotations in the digital world: users like to be able to post their thoughts, comments, and reactions to existing information. This form of *active presence* is important. Lombard concluded that users' sense of being in a space can be greatly enhanced by "the ability to interact with mediated environment [Lombard97]." A second benefit exists: facilitating active presence can also contribute to the accumulation of information.

Active presence, though potentially more rewarding for both the contributor and others, does require more effort from him. A good presence system should allow for both forms of presence. Passive presence provides a more inclusive overview of users' interests, while active presence enables them to contribute more thoughtfully.

2.6 Providing information richness without sacrificing privacy

Visualizing social presence means showing information about individuals. However, people may want their interactions to be kept private. We seem to be present with a dilemma. Total privacy would deprive everyone an opportunity to enjoy the benefits of presence we have discussed. Instead, we need to find compromise that would create a vibrant on-line space without disclosing sensitive information about individuals.

There are several ways to address this problem, some of which can be combined:

- Anonymity -- not identifying any individuals. Instead, show only collective information, such as the number of people at different parts of the space and the most popular paths.

- Pseudonymity -- identifying users by alternative names. This is useful for showing both overall patterns about an environment as well as individual paths, without disclosing personal identity. This scheme has been used in Judith Donath's Visual Who display [Donath95].

- User control -- allowing users to decide what information they want to release. People naturally want control over how they are presented to the public [Goffman59]. They should be able to decide which information is to be shown.

- Exchange of information -- giving away personal information in exchange for useful information. This follows the collaborative filtering model seen at sites like EachMovie [EachMovie98]. People are willing to list what movies they like, in return for recommendations of other movies they might like. Similarly, users may be willing to tell others about their interests in order to find others who share their interests.

- Community – sharing information within a community. Members of a pre-defined community may trust each other enough to share certain personal information. People from outside the community should be given less access to the information. To realize this, one needs means to create these communities and define their boundaries.

2.7 Summary

Social presence is largely missing from existing information spaces. This means that users cannot see the flow of on-line activities or the gathering of on-line crowds. As more people spend more time on the Web, the need for enhancing their Web experience by providing social presence becomes more urgent.

We have defined social presence and discussed its functions for social navigation and human connection. We then presented ways in which social presence can be supported. Current on-line environments need significant enhancement to reintroduce the social presence found in physical interactions. By implementing and extending some of the ideas here, we can create more humanistic environments that provide a much more enjoyable and enriching experience.

3. Interaction in the context of information

We have proposed that adding social presence to an information space can lead to interactions between individuals that enrich their on-line experience. Just how an interaction comes about is the topic of this chapter. While social behavior is an important part of sociology, little has been written on steps leading to an interaction. Instead, we have on one hand broad studies of sociology [Aronson54] or social behaviors [Blumer86], and on the other hand detailed line-by-line psychological analysis of an interaction [James71].

In his *Behavior in Public Places: Notes on the Social Organization of Gatherings*, Erving Goffman presented a useful framework [Goffman63]. He made careful observations about how people behave in public places where they can be observed or approached by others. Noting that social behavior is hard to examine in day-to-day life, where much of it seems intuitive, he examined in particular mental institutions, where the deviations from normal behaviors are the most obvious. He was able to explicate many of the small but important rules that govern social gatherings. We will start with Goffman's framework for examining these rules, then look at how face-to-face and on-line interactions differ. Finally, we will introduce our own framework of interaction ramp-up.

3.1 Goffman's framework

Goffman found two distinguishing features of behavior in public places: richness of information flow, and facilitation of feedback. In real life, a person is capable of perceiving a variety of cues about someone he encounters, such as posture, clothes, looks, facial expressions, gaze, speech, hand gestures, etc.. He is also subject to the same perception himself. He might therefore exercise certain control over these cues to be considered presentable or worthy of interaction with another human being. Different cultures have different sets of rules about what is acceptable behavior in various circumstances.

Goffman then established the following framework for examining the set of public behaviors, as shown in Figure 3.1. Let us examine this framework in more details.

```
┌─────────────────────────────────────────────────────┐
│                                                     │
│  - Unfocused (presentation, involvement)            │
│                                                     │
│  - Focused (civil inattention -> face engagement)   │
│       - Acquaintance                                │
│       - Unacquainted                                │
│              - Exposed position                     │
│              - Open position                        │
│              - Mutually open position               │
│                                                     │
└─────────────────────────────────────────────────────┘
```

Figure 3.1 Goffman's framework for behavior in public places

3.1.1 Unfocused Interaction

Unfocused interaction denotes public gatherings where there is no common focus of attention. A random group of people waiting at a bus stop is an example. *Focused interaction* denotes where there is a shared focus for possible interaction, such as when a few of the people start a conversation.

In the case of unfocused interaction, there are certain rules about how individuals are allowed to present themselves or involve themselves in activities in public. One should exercise certain control over his faculties, remain alert and open to the possibilities of interaction. In an American middle-class environment, it is unacceptable to be totally uninvolved in any activity for extended periods of time. Thus people may adopt tricks like taking a cigarette break or reading newspapers. On the other hand, it is also unacceptable to be totally involved in an activity without regards to others in the environment. In the years since the 60's when Goffman's book was written, our societies has certainly seen a relaxation of some of these rules about acceptable public behavior. New gadgets like portable game console, palmtop computers, and especially cellular phones, have created new means for people to become more intensely involved in their own affairs in public.

3.3.2 Focused interaction

In the case of focused interaction, there is a gradual build up from *civil inattention* to *face engagement*. Civil inattention means not staring directly at another person, but giving him the benefit of the doubt as a fellow

participant in public gatherings. When strangers are confined to close quarters, such as during an elevator ride, we can often find people studiously trying to follow the rule and avoid staring at each other.

Face engagement means two or more people coming into an interaction. This can be further divided into engagements among people who know each other or not. *Acquaintances* would need a reason not to engage in interaction: even if they are far apart in a crowded room, one should still acknowledge the other, by a nod or smile.

The *unacquainted*, however, need a reason for engagement. This kind of interaction is, of course, the more interesting aspect of public behavior for us. Some groups in our society are in what Goffman calls an exposed position, and can be engaged at will, such as the very young and the very old. Clerks in stores, policemen on the streets, or an amusement park worker in a Mickey Mouse suit, are also in an *exposed position* because of the role they take on. In our society, uniforms or badges are often used to denote these roles.

There are also certain acceptable *open positions* that people can take to approach others. We can freely engage people in exposed positions. This can sometimes lead into a face engagement with others. For example, talking to a child may lead into a conversation with the mom about the age, health, etc. of the baby. In our society we can ask anyone in public for the time, direction, or coin change. These are considered free favors: we can also be called upon to help others in the future. We see giving people more possible opening positions can lead to more interactions.

Finally, there are certain environments that produce *mutually open positions*. For example, gated environments like parties allow all present to safely engage others. The English sports field presents a similar example; the friendliness among the English sportsmen often surprise their supposedly more open American counterparts. Places with small number of people, such as the American frontier or many of today's small towns, allow the same freedom. What about more crowded and diverse environments? Interestingly, being of the same nationality in a foreign country, or being members of an underprivileged minority of a country can also produce the same mutual openness for face engagements. Presumably, in contrast to other participants, being of the same nationality or

ethnicity may provide a strong enough commonality. Some of these situations might be reproducible on-line, such as by showing more commonalties between people.

3.2 How on-line interaction is different

Having seen Goffman's framework for face-to-face interaction, let us examine how on-line interaction is different. First, the rich cues are missing. In typical on-line environments today, users can identify themselves only with a username, and sometimes one of a given set of avatars. This leads to surprising innovation as people find ways to express themselves, as well as masquerades: someone who calls oneself TheHotBlonde is likely to be a male pretending to be female. It can certainly be liberating to not be judged on the basis of physical features, but also confusing for others to figure out who is who.

Feedback is also limited in current on-line environments. Because on-line environments are mediated, someone who has not spoken in a while could be listening to others' conversation or could be busy attending to real life affairs. People are not subject to the same perception for their behaviors; they in turn have poorer feedback from others about their own behaviors. However, Lynn Cherny found in her study of on-line environments like Multi-User Dungeons (MUDs) that people still creatively try to give feedback, or what is often called back channeling [Cherny94]. MUDs are predominantly text-only environments. Back channeling in such an environment can consists of saying "uh-huh" or "yeah?" to indicate interest, or making a smiley face (☺) to indicate a joke.

These missing informational cues and feedback affect the utility of various parts of Goffman's framework. Many of the basic ideas behind Goffman's framework survive intact, with the specific rules relaxed. From our observation of several Web-based message boards and chat rooms, we observe that people often still feel the need for an opening in an on-line encounter among the unacquainted. We also often see that people who come back to an on-line conversation would apologize for real-life interruptions that took them away. We speculate that, like their counterparts in the real world, a gated or less popular Web site may produce greater openness among its participants.

Other parts of the framework , however, no longer apply. For example, individuals need not be properly dressed for an on-line occasion, nor have an acceptable level of involvement themselves. One can conduct several on-line text-

based conversations at once, if one's fingers can keep up. It also is perfectly okay to stare at someone's pictorial representation on-line if no video feedback is provided. Of course, a static picture is also not quite as interesting an object for one's stare.

3.3 Interaction Ramp-up

For a better framework for describing on-line interaction, we have created a set of stages for gradual interaction. We call this the interaction ramp-up.

1. Presence - seeing other users

2. Patterns - finding trends in where people go and who are going

3. Focus - identifying groups of interest and being able to monitor their behavior

4. Opening - expressing interest in conversation

5. Conversing - response to officially start a conversation

Figure 3.2 The interaction ramp-up.

Some of these correspond to different parts of Goffman's framework, as we will show in more detailed discussion below. Also, some of these stages can be skipped under certain circumstances: when we see someone we know on-line, we can immediately skip from presence to focus and open a conversation.

3.3.1 Presence

We have discussed social presence in the previous chapter. As we mentioned, on-line environments often display current presence using usernames or avatars. In many chat rooms, an avatar simply indicates someone being present. We can often see the same people wondering around chat rooms in search of other people or happenings. In a more structured informational space, one's presence may indicate an interest in the information presented. For example, two people both visiting a Web page about Paris may share an interest in the country. We are interested in this more informative sense of social presence, which has a higher potential of leading to successful interactions.

3.3.2 Patterns

Because movements in these structured environments convey more meaning, a new participant or a returning participant can observe what others do and to gain useful insights about other participants:

- Where do people tend to come from or go to?

- When are there more people?

- What can we learn about their interest?

Answering these questions can help one find relevant other places to see or people with similar interests to interact with. Different on-line environments offer varying amount of support for determining these patterns. In Usenet newsgroups, a globally distributed interaction forum, we can tell quite a bit about a poster's interests based on what he or she posts, but very little about someone who only reads. For-sale postings, for example, might be read by many and yet receive no on-line replies. It is difficult to assess the amount of interest in them. For Web pages, counters can indicate the accumulated amount of interest about a page, but does not tell much about current usage. Slow downloading speed may indicate a crowd of current users, but may also signal a bad network connection. It is an imprecise, or in any case less desirable, indication of current accesses.

The presence and patterns stages correspond to Goffman's unfocused interaction. There is no shared focus of attention. Individuals are simply presenting themselves in a space and observing others also in the space. From here, one can follow others to find popular information, or move on to later stages of social interaction. Goffman's notion of involvement for unfocused interaction still applies: people on-line often need other focus of attention besides interaction. In a chat room, someone who is not actively participating in the discussion can easily become bored: one can only listen for so long, especially if the conversation is not of shared interest. On-line environments with more creative outlets can be quite attractive. For example, AlphaWorld enables users to build virtual worlds in addition to interacting with others [AlphaWorld99]. This can also lead to interesting conversations as users discuss what is being built. Being able to interact while one browse the content of the Web can also have similar effects.

3.3.3 Focus

Focus means identifying potential interaction partners, or identifying a topic of interest to pursue. In many chat rooms, one can simply start with a simple hello to the group at large. Depending on the friendliness of the environment, one may or may not get a response. A more focused greeting, either directed at a person or asking a specific question, would obtain a much better response.

As in the last stage, the ability of a user to determine focus depends much on the amount and kind of information available. We can certainly sort through a large quantity of chat messages to identify whether someone is a potential interaction partner. User profiles can also be informative. But sorting through text messages or user profiles can be time-consuming. Keyword searches may or may not help depending on whether one knows the words to search for. We see the usefulness of uniforms in real life to quickly establish information about someone.

Virtual environments need to provide more support for its users in determining whom to talk to. Here the granularity of information might make a difference. If we know someone has read not only one but several news stories on the same topic we are also interested in, we can more safely assume he has a similar interest. However, this information is currently not readily available at most Web sites.

3.3.4 Opening

Being able to identify potential interaction partners leads nicely into the Opening Stage. This stage is where one indicates interest in talking to someone, by a greeting or question. Here Goffman's exposed, open, or mutually open positions aptly apply. Moderators, chat hosts, or shop assistants in a shopping site are in exposed positions. One can talk to chat bots or people with funny names that suggest they are more open for interaction. An example of an open position is asking people how to use a particular chat feature.

More interesting is the ability to talk to others sharing similar interest, which can lead to more fulfilling conversations. Here, we must not under-estimate the importance of icebreaking. The Meme Tag was designed to help people find others with shared interests: when two people who profess similar interests come near each other, the devices on both would blink. But to the designers' surprise, they found that it was more useful as a simple icebreaker between strangers who otherwise may not feel comfortable approaching one another [Borovoy98].

People may need various support for opening up. Jane Jacobs gave an example of a local shop owner who never directly introduces his customers who he knows to share a certain interest, because he respects their privacy. Instead, he would start talking about the subject of interest, and leave the customers to decide whether they want to

45

pursue the topic [Jacobs61]. We may want to create virtual shop owners that can both respect users' privacy and lead them naturally into conversation.

One type of possible open positions are comments on subjects of mutual interest to whomever is present. Real-life etiquette books suggest that we can always talk about the weather, apparently relevant to everyone. We are also told to follow current events so that we can have interesting things to say, along the line of "Has anyone heard about...?" There are also more appropriate opening positions based on ones' surroundings or time. At a school for example, talking about upcoming school events is relevant. During Halloween, talking about costumes is appropriate. We see that these ideas have natural counterpart in on-line environments. Different comments are relevant at different Web sites. Helping users to identify these relevant topics can lead to better openings and more focused discussion.

3.3.5 Conversing

The last of the five stages is when someone responds to an opening move and enters into a conversation. The more two users share interests, the more likely the conversation will continue. Because on-line environments are mediated, individuals have less embarrassment in being rejected for a conversation. Thus we can see the usual junk openings and lack of responses in many chat rooms. One way to measure the successfulness of discussion boards is by the response rate. By creating an interaction environment that allow users to move through the five stages step by step, first understanding the interaction environment and then identifying potential interaction partners, we increase the possibility of a successful interaction.

3.4 Summary

Because the rich informational channels are missing, and the feedback is not as strong, on-line interaction poses certain unique challenges. We have created the interaction ramp-up, a five-stage framework for describing on-line interaction, and compared it with Goffman's hierarchical framework for examining face to face behaviors in real life.

A framework is only as useful as the questions it answers or brings up. We will see how our framework can translate into useful design requirements for the designs in the following chapters.

Part II. Experimental designs

We will explore the design space for visualizing information structure and social presence through several designs for the Web and the Usenet, a globally distributed discussion forum. The three designs are smaller pieces that each explores a different but related idea: from how user presence leaves its mark an information space to examining interconnection between people and information objects (such as postings), and then to more close-up and detailed portraits of individual users. The table below shows how the three designs differ in their emphasis on the display of information structure or people. We will distill useful ideas from these designs, ideas that will be incorporated into the larger real-time system described in Part III.

	WebFan	NetViz	PeopleFlower
Visualizing information structure	✔	✔	
Visualizing users		✔	✔

Where the three designs fall in the design space.

Visualization

To understand the structure of the information space and the related social data, users need to be able to quickly apprehend a large amount of information. Since people are more adept at processing complex visual relationships than lengthy textual descriptions [Arnheim69], our solution is to visualize information content and social presence as the basis for interaction support. Visualization allows a large amount of information to be conveyed compactly to the users [Card99]. Our work emphasizes visualization for average users:

- Legibility - helping users to grasp easily the information structure and social presence

- Aesthetics - creating interesting and attractive visualizations that can be a focus for users both in generating new conversation and a mirror for the user community to reflect on its activities

- Exploration - allowing users to interact with the visualization and select areas of interest for more information.

4. WebFan: Visualizing presence on the Web

We have discussed the benefits of providing presence in an information space and discussed how it can lead up to interaction. WebFan, the first of our design pieces, explores how to visualize the activity of users at a Web site. We are interested in helping users with the first three stages of the interaction ramp-up: presence, patterns and designs.

In Chapter 2, we discussed how to show past presence of digital objects by showing accumulated usage, and how to represent the current presence using histograms or individual avatars. Here, we are interested in providing both past and current presence to help people identify patterns and individuals of interest. Our approach is two-fold: first we visualize the structure of information space, then overlay real time presence information.

4.1 Previous Solutions

As mentioned in Chapter 2, Footprints visualizes historical presence. The map view shows how various pages of a Web site are linked based on where Web visitors tend to go; the path view shows from a single given page, where are the other popular destinations. This visualization scheme is appropriate since the goal for the project is social navigation: using others' traversals to help navigate a space.

Closer to our work are projects that address live user presence and interaction on the Web. One of the earliest works in this area is the Sociable Web[Donath94]. It consists of a modified Web browser and server system that shows which users are on the same Web page, and allows the user to strike up a new conversation or join an ongoing discussion. Users can see a textual listing or a pictorial view of other users. More recently, Nelson Minar's crowd visualization shows a loosely clustered set of pages and their visitors to give a sense of crowds at a site [Minar99]. The dots
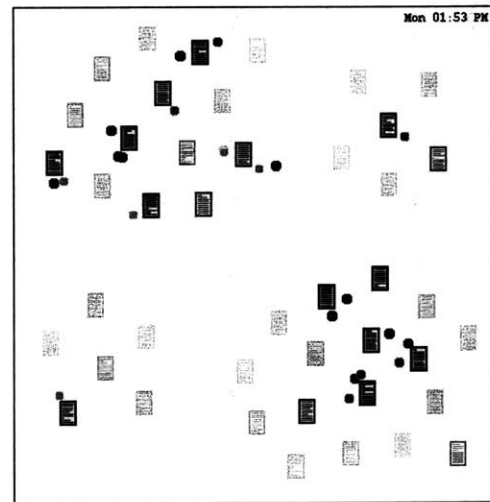

Figure 4.1 Nelson Minar's Web crowd visualization shows people as dots moving around rectangular pages [Minar99].

representing people jiggle slightly to convey effectively a sense of people milling around a Web page. These works

have not dealt with visualizing activities in a structured Web site, where users' movements can help people identify potential interaction partners. In terms of the interaction ramp-up, these previous methods offer support for the presence stage, but only limited (if any) support for the following stages of patterns and focus.

4.2 WebFan's approach

WebFan visualizes both past and present traversals based on a hierarchical representation of the Web. The hierarchical structure of the Web site can be created by hand or by running a Web crawler through the Web site starting from the main page.

We will focus our discussion on Web-based message boards (Web boards), which contain threaded messages that each resides on a separate Web page. These boards have a natural hierarchy based on the threading of conversation, which make them well-suited for our visualization. Web boards are also interesting since the content is itself conversations, although in this case asynchronous. In addition, they provide a challenge in that their contents change over time. Although we will focus our discussion on Web boards, WebFan can be used to visualize any Web site with hierarchical structure. A similar tree visualization has been used for comparing temporal snapshots of Web sites to help site administrator maintain their sites [Chi98].

WebFan uses the thread structure to lay out the messages in a fan-like fashion, as shown in Figure 4.2. Each line segment represents a message, and segments branching from a segment represent subsequent replies. The most recently started threads are shown at the top, and older threads follow in a clockwise order. The tree-like layout is used because the underlying structure of the content is hierarchical. A semicircle outside the tree structure encloses the tree structure, giving a sense of wholeness. I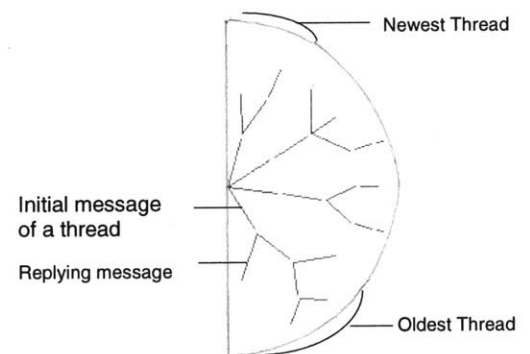t is also used to show angular displacement of the cursor to aid in navigation. Here a semicircle is used rather than a full circle to give a definite start and end to the visualization for showing the relative age of

Figure 4.2. WebFan layout

50

threads. This shape also makes navigation easier, as developers of the inter-cite clan graph, a semi-circular layout of related Web sites, found out from their users [Terveen98].

4.3 Visualizing site structure

This simple visualization can be used for users to quickly gain an overview of a Web board. The number of messages displayed show whether it is a busy board; how far the threads extend radially shows whether the discussion gets more involved. One can visually observe changes over time or identify conversational structure that stands out.
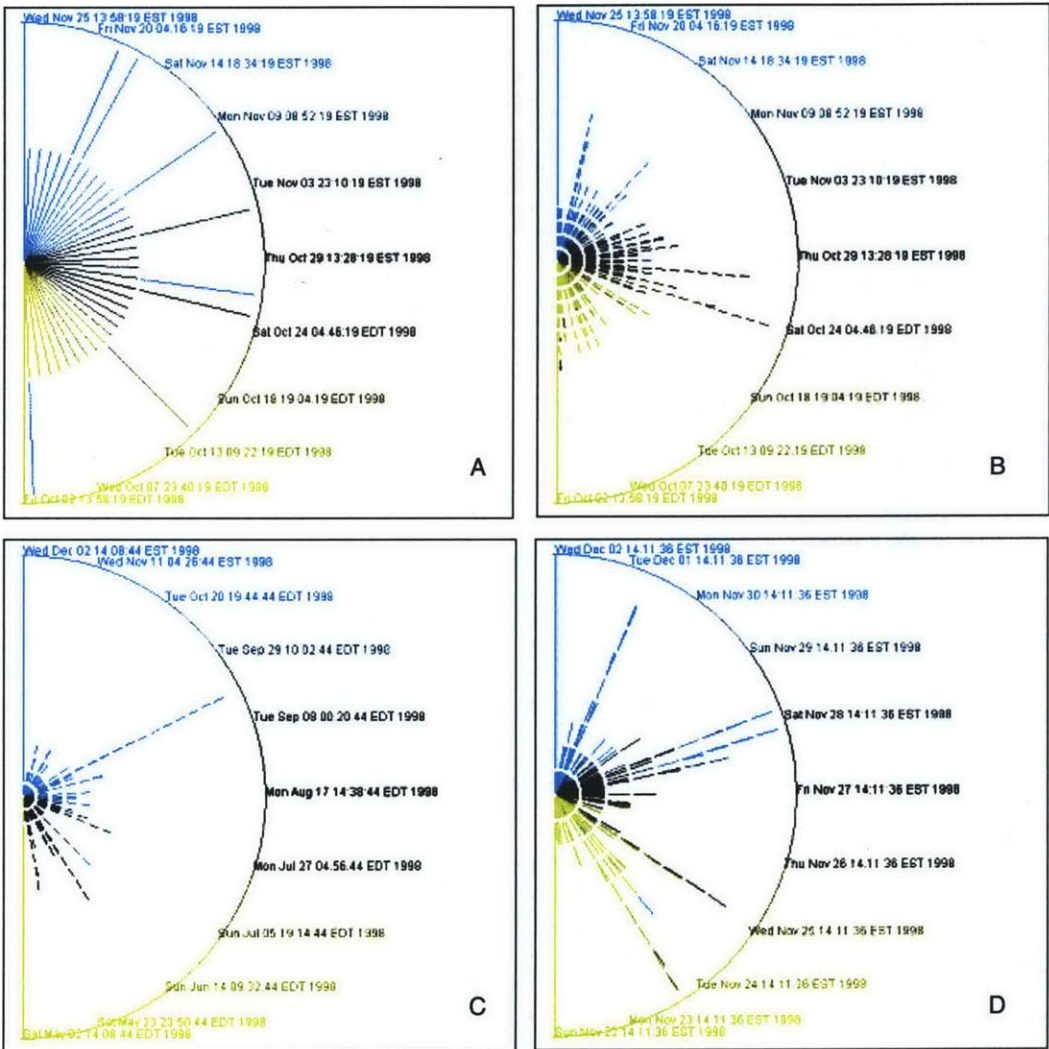

Figure 4.3 Four WebfFans illustrate the difference between four Web boards. Coloring shows when messages are created.

Figure 4.3 shows the WebFan visualizations of four Web boards. They show at a glance that the structures of the boards differ. Board A is a for-sale board, so most messages have no responses or only one response. The other boards all involve more discussions, so the depth of the threads is higher. Board B is the busiest board of the four, with both more posts and deeper levels of conversation. We can also see anomalies like the extra-long thread one third from the beginning in Board C. It is a topic that provoked much debate with in the group.

We can further use different visual parameters to illustrate information about the messages. In Figure 4.3, the color of the messages represent time of posting: from recent (cyan) to medium (gray) to old (yellow). The text around the semicircle gives a more detailed legend about what color corresponds to when. We see that overall, people tend to post and respond in about the same period, so adjacent messages are similar in color. But in Board D, we see several cyan messages in older gray threads or gray messages in older yellow threads. In this board, people tend to respond more to older messages. This could be a board where people check less often and react more slowly to postings. The coloring could also be useful to help users identify newly posted messages to older and therefore less noticeable threads. This is task is much harder to perform with a textual view, where users cannot easily gain an overview of the messages. Finally, we note that coloring the message by time also shows an anomaly: Board C has very few yellow threads. This could mean that the board initially was not very busy, but after a while picked up much more interest.

A busy board can quickly accumulate a large number of postings. The WebFan in Figure 4.3 seem rather crowded as a result. One useful technique for dealing with large amount of data is to show both focus and context. Figure 4.4 shows a WebFan expanding to show more details based on the position of the cursor. Users need to see clearly both what they are currently focusing on and how it fits into the greater whole. Lamping's work on hyperbolic tree visualization is based on this important assumption [Lamping95].
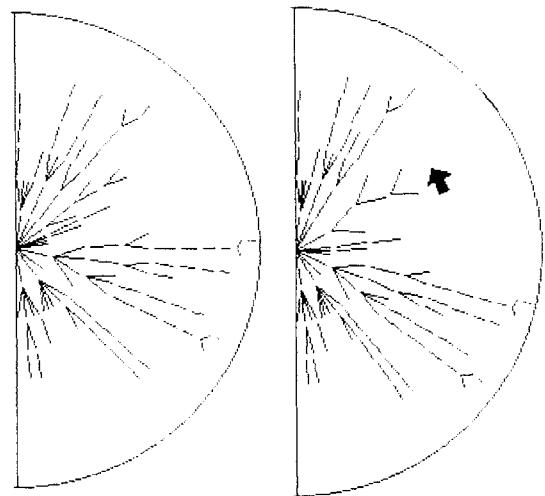

Figure 4.4 Two views of the same WebFan showing how the fan expands to show more detail around the cursor.

For spatial navigation, it is particularly important to see the context for a user to not get lost. We have made the expansion around the cursor smooth, so users can easily move around to examine various parts of the fan. We have borrow the metaphor of a Rolodex: as a user flips through the threads, different part of the fan expands and compresses in turn.

4.4 Visualizing user activities

We discussed how WebFan can be colored to show additional information about the messages, such as the time of posting. We can use a different color scheme to dynamically display user activities. We show all accessing and posting activities within a selected time period. The visualization can answer user questions about overall patterns and specifics such as

a)   Which messages are other users viewing at any given instant? (Where can I find other people?)

b)   Which messages are being and have been read by many people? (What is hot?)

c)   Which users have read a particular message? (Who is interested in what I am interested in?)

d)   What messages has a particular user been looking at? (What is this person interested in?)

We are interested in helping participants find out about each other, so individual accesses are shown, unlike Footprints and similar works where only cumulative amount of past history is shown.

Each user is assigned a different color. Each time the user accesses a message residing on a Web page, the initially gray segment representing the message will change to that color. If multiple users have accessed a message, the segment corresponding to the message will have multiple colored sections representing the different users. A circle on the colored section indicates that the user is currently viewing the message.

Figure 4.5 is a Web Fan showing the access data of a Web board at two different hours of the day. The left view shows access during the early part of the day when the traffic is light. The right view is from later in the day when more traffic has occurred. One can see that the newer messages near the top are accessed more often, as indicated by their coloring. Also, due to posting of new messages, the shapes of the top few threads have changed.
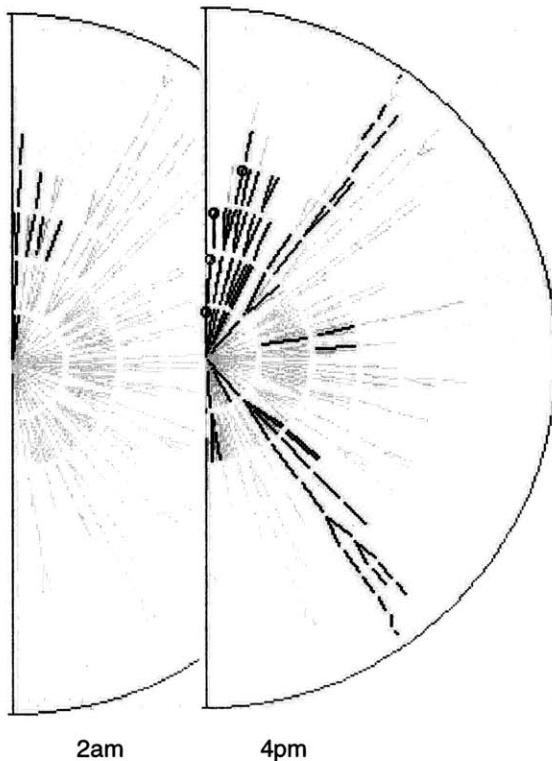
53

2am   4pm

Figure 4.5 Two views of a WebFan showing
accumulated access over a day.

We can use WebFan to easily answer the earlier questions:

a)   The number and location of circles show how many users are currently reading messages, and which messages they are reading. Here we see four current users near the top of the fan reading the latest messages.

b)   A segment with many colored sections represents a message read by many users. We see quite a few examples near the top of the fan. Not surprisingly, the latest messages are popular among current users as well as past users.

c)   The colors of a line segment indicate which users have visited the corresponding message. This is useful for finding others who have read a particularly interesting message.

d) All messages viewed by a user have been marked by that user's color. For example, towards the bottom of the fan, we see that an entire thread has been colored in blue, indicating someone very interested in this topic. If a new user share this interest, this could be a good start for a conversation.

4.5 Temporal View

WebFan shows a view of user presence in the information. Users may have more questions about the temporal aspect of user presence:

e)   When are there more people at the board? (When is a good time to look for others to interact with?)

f)   How long does a particular user stay around? (Does this guy find the board interesting?)

To answer e), a simple histogram representing the number of users at different times is sufficient. However, to answer f) requires knowledge about each users' stay at the board. We use small color bars to represent each user session at the board; the length of the bar indicates length of the session. The naïve layout, using a new column for

each user, wastes valuable screen space. Figure 4.6 shows three possible alternatives. Method 1 wraps the bars around at fixed number of columns; Method 2 also wraps the bars, but only if the current column is taken; Method 3 looks for the last open slot going left from the current slot, in effect compacting the bars to the left.

We see that because the first two methods both start from the current slot and add bars to the right, they create a misleading diagonal pattern with no real meaning. Method 2 is slightly more informative because a vertical stretch of blocks represent a period where there is only one user. However, this is not extensible to busier sites where such periods are rare.

The last method is by far the best, because the right outline of the bars creates the effect of a histogram. One can tell how many visitors have been to the board at a glance: 1-2 for the most part, 3 occasionally, and 7 during a particularly busy time around 4pm. Most people tend to visit the board for less than half an hour; a few stayed on for an hour and more. Note that the histogram created is a slight over-approximation. For example, if one user enters right before another leaves, two columns are occupied when there is only one user (circled in figure 4.6 for Method 3).
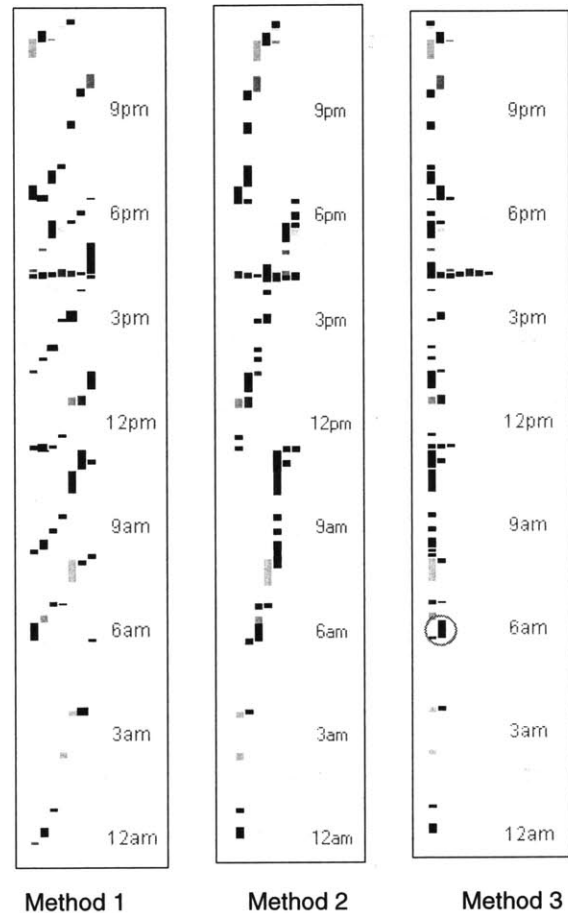
Method 1   Method 2   Method 3

Figure 4.6 Three different methods for showing how long the different visitors to the board stayed.

The WebFan and the temporal view are placed side by side as Web presence is displayed, as shown in Figure 4.7. Our goal is to allow users to obtain a quick overview of the data. This way, user can leave the visualization up while focusing on other things, using WebFan as a background visualization to gain a sense of other users' presence. When more activities are taking place, users can then refocus on the visualization. They can explore it in more detail to narrow down on individuals of interest. As one mouses over the temporal view, one can highlight a bar, and see all

the messages read by the selected user highlighted on the WebFan. This way, one can tell both who are around, and what messages they find intersting.
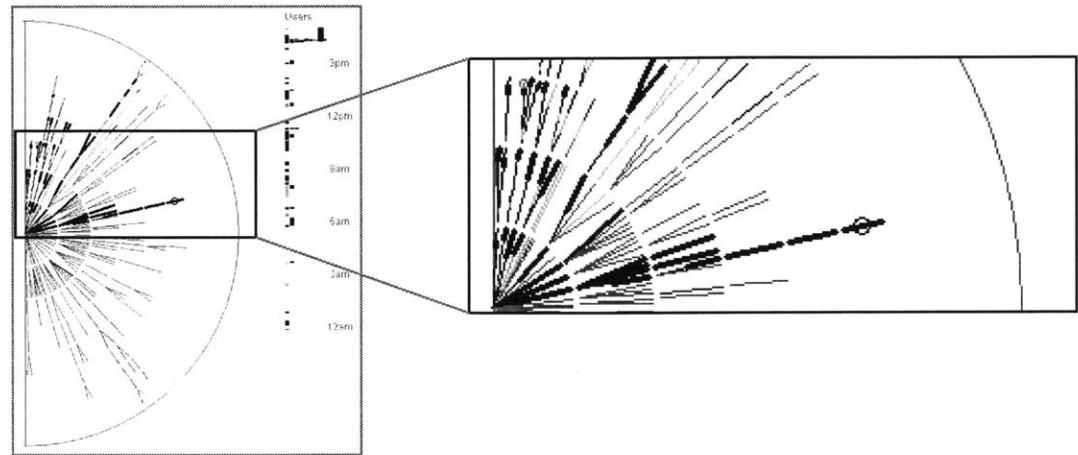

Figure 4.7 An enlarged view of a WebFan highlighting all the messages viewd by a particular user (colored black). Note that this user is currently at the board, as indicated by the circle on the rightmost message.

4.6 Discussion

We see WebFan as a useful supplement to the current Web browsing interface. WebFan has been implemented as a Java applet. Users can view the activity flow in the applet in one frame, and select a message to read its detailed text in another frame. Both recorded Web accesses and real-time access can be used for the user presence animation.

We would like to apply WebFan to generic Web sites by first extracting the hierarchical structure of a Web site then visualizing its accesses. Note that WebFan assumes a hierarchy for the pages, which works well for Web boards with such a structure. We will see in Chapter 7 how to visualize sites without a strict hierarchy.

On a Web site with more users, scaling issues will need to be addressed. There are really two parts to scaling: sites with more pages and sites with more people. For larger sites, it can be hard to see all the pages at different levels even with the Rolodex expansion. How to solve this problem depends on the nature of the sites. For sites that can be easily divided to different unrelated parts, we can use different WebFans to represent different parts of the Web site. For sites where the first few levels are most revealing, we can limit the visualization to the top few levels, or have several levels of visualization, embedding fans within another.

For sites with more visitors, one way is to allow the viewer to specify Web visitors that they are interested in, such as by domain or by places visited. The other method is to group together visitors of similar interest or background. Color can be assigned to different groups rather than individuals, but with the side effect that we will lose a sense of individual paths. With more information about the visitors, from user profiles, for example, such filtering and grouping tasks can be simplified.

Other Web board users' behavior can be fascinating to watch and follow, much as people-watching in a crowded mall or on a boulevard. Whereas real-life activity is often transient, activities on the Web can be recorded, searched, and played back. How much information about a user should we display? WebFan does not show the IP address of users, only their paths. But as we gain gather more information about the visitors over time, there is a real temptation to reveal more about these people and create a more informative scene. Here we see the importance of balancing information richness and user privacy, as discussed in Chapter 2.

4.7 Design lessons learned

We have presented the goal and design of WebFan for visualizing user activities on a Web Board. WebFan explores how to support the initial three stages of interaction ramp-up: 1) presence, 2) patterns, and 3) focus. It can show presence in both the content view and the temporal view. We have also discussed how it can be used to determine overall patterns of activities at the board and individual behaviors.

How successful was WebFan in addressing these questions? We have received some informal feedback. One problem we find is that the visualization is still somewhat complex. Many users do not have a visual image of the conversational structure at a board, so the tree structure is less intuitive for them. It takes them some time to understand the visualization and be able to interact with it.

We also received many positive comments. Users are receptive to the idea of a compact visualization being used as an interface to large amounts of data. Two things people liked the most about the WebFan are 1) being able to play with the interface and see it expands in response and 2) watching the dynamic animation of user activities. A few

57

commented favorably on the aesthetics: WebFan consists of simple lines that are arranged in a meaningful and visually pleasing manner. On the more practical side, people liked being able to zoom in on what appears to be an interesting thread and read the actual messages. This makes the visualization much more useful as both a navigational aid and a representation of activities at the board.

Some people had problems selecting a particular message. This could be because the user needs to move both radially and angularly, which is not easy at first. We spent quite a large amount of time trying to get the Rolodex effect correct to help them better see where they are and navigate accordingly. Initially, only the thread containing the selected message is expanded. This makes selecting messages within a thread easy, but going from thread to thread difficult. Users often lose their place. So we made the expansion more smooth, and based on angular distance from the position of the cursor. This makes moving from thread to thread easy, but within thread difficult, since the visual expansion is too responsive to small movements within a thread. Finally, we kept the smooth expansion, and allowed users to lock down the expansion when they have selected a thread of interest. This facilitates both kinds of movements.

Another lesson we have learned is that selecting the right initial design is important. Often the simpler the initial design, the more easily it can evolve into an effective final design. Once we decided on the basic fan visualization, we spent quite a bit of effort to make it easy to use. We are happy with the information presented by WebFan, and its interactive and dynamic aspects. However, we feel that if a simpler design is chosen in the first place, we could have saved much more effort later. Perhaps this lesson can only be fully internalized through actual designs.

A quick summary of what makes a good design

- Simple initial design
- Dynamic flow
- Interactive exploration
- Aesthetics

A demo and some screenshots of WebFan can be found at:

http://graphics.lcs.mit.edu/~becca/webfan/

5. NetViz: Visualizing connections between people and objects

The previous design, WebFan, visually represent data on which pages are viewed by whom to help users identify interesting postings and potential interaction partners. NetViz, our next piece, explores further the interconnection between people and informational objects. I created NetViz with Marc Smith and Steven Drucker during a summer internship at Microsoft Research.

NetViz visualizations, like the ones in Figure 5.1, show how different Usenet newsgroups are related, as well as how topics and posters within a newsgroup are related. This helps users to see overall patterns, identify potential areas/posters of interest, gain knowledge about the interaction context, and interact more effectively. NetViz is based on a framework that deals with how objects and people are connected. We start by explaining this framework.
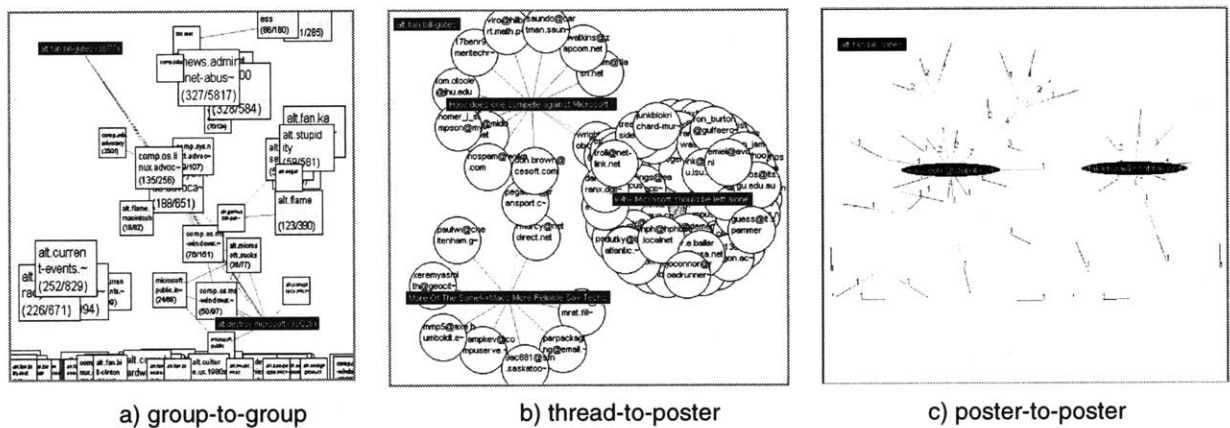

| a) group-to-group | b) thread-to-poster | c) poster-to-poster |

Figure 5.1 NetViz shows how groups, threads and people are connected

5.1 Framework of interconnections between people and objects

How can we clarify the interaction context of an on-line environment? In environments like Usenet newsgroups or the Web, information about the kind of space, group and activity taking place, is usually missing or ambiguous. An informed participant needs to be aware of the underlying network of relationships between people and objects, such as topics of discussion or Web pages. We have created a framework for thinking about such interconnections. We will present it using a simple example about the relationship between book buyers and the books they purchased.

BOOK BUYER	BOOKS PURCHASED
PersonA	Book1, Book2, Book3
PersonB	Book1, Book2, Book3
PersonC	Book1, Book4
PersonD	Book4, Book5
PersonE	Book4, Book5

Table 5.1 Book purchase information.

Book purchase information, as shown in Table 5.1, is commonly used in collaborative filtering [Goldberg92, Resnick94] to recommend books to users of shared interests [Amazon99]. Here, we are interested in not just finding out about books but also about users themselves. We would like to answer user questions such as

- Which books are related?

- My taste in books is similar to X's. What are the books I might be interested in?

- If I like a particular book, who else shares my interest?

- Which users have similar general tastes in books?

5.2 Four diagrams of the interconnections

Trying to answer these questions using data in tables can be difficult, especially when we have a large number of books and purchasers. We can look at these relationships instead through four related views:

- *Object-to-Object* – displays connections between objects based on user activity

- *Object-to-Person* – starting from objects, see who are the related people

- *Person-to-Object* – starting from people, see which are the related objects

- *Person-to-Person* – how people relate to each other based on interest in similar books

Figures 5.2-5.5 show examples of the above. In these figures, rectangular nodes represent objects (in this case books) and circular nodes represent people. The lines connecting them represent various relationships.

5.2.1 Object-to-Object Diagram: showing relationships among objects

The Object-to-Object Diagram shows the relationship among objects based on user preferences, as shown in Figure 5.2. The nodes represent books, and the links represent how often the same user purchases both books (the more
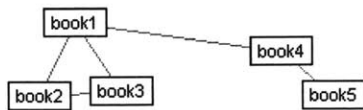
60

Figure 5.2 Object-to-Object Diagram shows which books might be similar based on user purchases.

often, the stronger the connection, and the shorter the link). By looking at the diagram, one can infer that book1 and book2 are more closely related to each other than to book5. In fact, book1, book2, and book3 form a cluster because they are all connected to each other. They may represent a set of books on similar topics. This diagram can help users see how a set of books may be related to each other and find other books on similar topics based on books they already know. A partial view of this information can be obtained in the form of book recommendation from Amazon today.

5.2.2 Object-to-Person Diagram: going from known objects to people

The Object-to-Person Diagram in Figure 5.3 reveals who are the actual people connecting the books in the last diagram. The rectangular nodes represent selected books; they are linked to circular nodes representing people who have purchased these books. Note not all the books are shown in this view: part of what makes
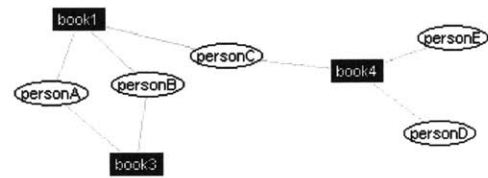

Figure 5.3 The Object-to-Person Diagram shows for selected books, who are people of interest.

this view useful is that the user can select only the books they deem relevant. By looking at the diagram, one can infer that personA and personB are both interested in book3, and may share interests in a similar set of books, whereas personE for example is interested in somewhat different books. If a user is also interested in book1, then the user might share interests similar to personA, personB, and person C. This view can help people seek out others with similar tastes in books.

5.2.3 Person-to-Object Diagram: going from known people to objects

We can also go the opposite way, from know people to objects, in the Person-to-Object Diagram. For a given set of people, say personA, personB, and personC, the diagram can display how different objects connect them to one another, as shown in Figure 5.4. The black nodes represent selected people and the links connect them to the books they have purchased. By looking at the diagram, one can infer that book1, book2, or book3 is closer to personA's interest than book4. So if a user's interest is more aligned with personA's, he might want to look at the first 3 books. If a user knows a group of people who share her interest, and another group who does not share his interest, The Person-to-Object Diagram can help separate out the books based on the preference of people whom he knows.
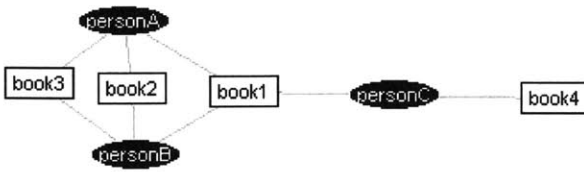
61

Figure 5.4 The Person-to-Object Diagram shows, based on selected people, which books might be of interest.

The Object-to-Person Diagram in the previous section also shows both people and objects. The difference is that the former starts out with known books and shows how people are related, and the latter starts with known people and shows relevant books.

5.2.4 Person-to-Person Diagram: showing relationships among people

Finally, we are interested in showing the relationship between people themselves based on their preference for objects, in the Person-to-Person Diagram. In Figure 5.5 for example, the nodes represent people, and the links represent how often each person purchases the same books (the more often, the stronger the connection, and the shorter the link). So by looking at the diagram, one can infer that personA, personB are more closely related to each other than to others. In fact, personA and personB or personD and personE may be a pair of people who share similar book purchases, and may share a lot of interests in common. PersonC, on the other hand, seems to share diverse interests with a lot of people.
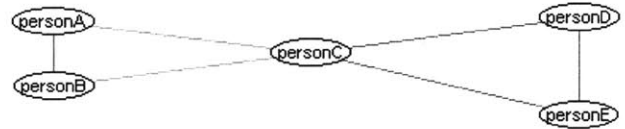

Figure 5.5 The Person-to-Person Diagram shows who shares similar tastes in books.

Notice the parallel and dual relationships between the diagrams. Both to-Object Diagrams both help users to locate objects, while the to-Person Diagrams help locate other people to interact with. Object-to-Object and Person-to-Person both show a single type of nodes. The former gives an overview of how objects are related, while the latter deals with people. The other two diagrams both show a mixture of two types of nodes. They use known relationships between one type of nodes to separate out the other type of nodes. Table 2 summarizes the similarities:

	Locating	Type of nodes
Object-to-Object	Objects	Single
Person-to-Object	Objects	Mixed
Object-to-Person	People	Mixed
Person-to-Person	People	Single

Table 5.2 A summary of the 4 diagrams showing what they can be used for and what nodes they contain.

So far, we have presented the four views in the most basic format. In the next section, we will look at how some of these four views have been used independently in previous works. The four views we have presented can be generated by existing techniques. We have created interaction techniques to reduce their complexity and enhance their usability. Finally, we will apply our visualization framework and techniques to the Usenet to reveal patterns and connections that are difficult or impossible to see in existing interfaces.

5.3. Related works

Our work is influenced by recent developments in collaborative filtering, as well as works in more established fields such as social network analysis and information visualization.

Collaborative filtering originated from the Information Tapestry Project at Xerox PARC [Goldberg92] for e-mail filtering, and has been applied to other domains such as Usenet postings [Resnick94]. By using recommendations from other Usenet readers, for example, a program can make recommendation of postings for new users. Our visualization can be used to for visual collaborative filtering.

Information visualization techniques are typically used for the visualization of large amounts of commercial, financial, scientific and technical data for data analysis and mining. Their main audience has been data analysts for examining data from different angles. It has also been used to help end-users filter data to select movies or find houses using dynamic queries [Ahlberg94]. Here we follow in the latter tradition in using visualization as a simple interface to help the average user.

Our work makes extensive use of networks, or directed and undirected graphs. Networks are used extensively for a narrow range of applications, most notably in telecommunications management and social network analysis. Unlike trees, networks may contain cycles, which are important for showing complex, non-hierarchical information. To limit the amount of information shown at once, and give users control over what they see, we have also used traditional information visualization methods like filtering [Mukherjea95].

5.3.1 Previous examples of four views

Let us see how the four views presented here have been used independently in previous work. Our Object-to-Object view can be seen as a new way to visualize collaborative filtering information, which shows how close objects are to each other based on user preferences. Alternatively, our work can be used as way to perform visual collaborative filtering, by directly looking at other people's preferences through the Person-to-Object Diagram. But more importantly, our visualization can show how the objects recommended fit into the greater web of connections among the objects and people, giving our users greater understanding and control.

Loren Terveen and Will Hill's Auditorium Visualization for Inter-site Clan Graphs displays the relative importance of different Web sites based on a user's query [Terveen98]. Web sites are laid out in a series of concentric semicircles based on importance. They used dynamic ordering of Web sites to allow user customization based on interest. This is similar in spirit to our Object-to-Object view with a weighted display, as described later.

Judith Donath's "Visual Who" displays the interconnections between people based on mailing list membership [Donath95]. A set of mailing lists are selected as anchors, the names of people are then placed around the anchors based on how similar they are to the members of these lists. Dynamic animation of changes in information and user interaction made the display easy to use. The visualization is similar to our Object-to-Person Diagram, which lays out people based on closeness to objects.

Henry Kautz's ReferralWeb, which also uses networks, shows a Person-to-Person Diagram where the objects linking the people are jointly published papers [Kautz98]. It is an example of a social network diagram. Social network analysis is the study of the structure and dynamics of relationships between people. Originating with Jacob Moreno's sociometric research [Moreno34] in the 1920's, social network analysis has demonstrated the dense interconnections between people that allow any two people on Earth to connect through six or fewer intermediaries [Milgram67]. Network analysis has also been used to illuminate the empirical structure of local communities [Wellman79] and inter-office patterns of communication [Krebs96].

5.4 Visual interaction techniques

The four views of collaborative information we have described can be created using standard force-based models. Because a large amount of nodes may need to be displayed, we provide other ways to filter the data, differentiate better among the different nodes, and select views of interest to the individual user. On top of the simple visualization, we have added dynamic retrieval and display, interactive user-control, and overlay of information on demand. Examples of the use of these techniques will be shown in Section 5.

Filtering and the falling metaphor

Not all edges are created equal. A slider can be used to filter out edges with weaker tie strength. The original visualization contains all the ties relating the nodes. As the user moves the slider to the right, edges with weaker strengths are filtered out until only those with the strongest strength remain. The nodes that are no longer connected to other nodes fall to the bottom of the display. Moving the slider back to the left will reconnect and bring back the nodes. This dynamic motion is used to help people to identify easily which nodes are being eliminated. Users can also tell at a glance which nodes are not connected: they are at the bottom.

Displaying structure

A given diagram may contain a large number of nodes and edges. To better see the structure of the interconnections, users can select the structured view, which displays a simple square or circle representation for each node. This can reduce the clutter on the screen and make the network structure easier to see.

Sorting

The nodes can be placed using the original spring-based model or by sorting. The latter places nodes with more connections (i.e. higher degree) on top of the display. It is useful for identifying what are often more important nodes.

Overlaying information

On top of the basic view, we can overlay other information. Color and size of the nodes can denote different information about a person or object. For example, we can use the color of the book node to represent its rating. This can help users identify at a glance who purchases books of what rating.

Linking views

The various views can be linked together for users to easily sort through the information. From an Object-to-Object Diagram, for example, users can select interesting objects then switch to an Object-to-Person Diagram to locate people who are related to these Objects. They can then switch to a Person-to-Person view to see these people in the context of others.

5.5. NetViz -- applying the framework to Usenet

As a real-life illustration of our approach we present NetViz, a system for visualizing the relationships between groups, posters, and threads in the Usenet. Usenet is a globally distributed discussion space that is often seen as an example of the tragedy of the commons [Hardin68, Kollock96], a public resource that has been ruined by overuse and pollution. While Usenet has not been consumed by its tragic qualities [Smith97, Whittaker98], better tools for navigating and visualizing this space may allow people to return and restore this public commons.

To create visualizations, NetViz uses data generated by the Netscan Usenet analysis system, located at http://netscan.research.microsoft.com. Netscan examines an ongoing stream of Usenet messages and extracts from each message's header relevant information such as its poster and subject. Netscan then builds a number of databases that organize this data, which can be queried to generate a range of measures of activity and relationships in selected newsgroups.

5.5.1 User questions

NetViz shows connections between newsgroups and among topics and participants within newsgroups. Our goal is to support users in the process of gaining general awareness about a newsgroup, by determining which other groups are similar to it and the range of opinions expressed in the group. We are interested in answering questions in three broad categories, about groups, threads, and people:

1. What groups are related to those that I am interested in? Which group has the most activity? How about the kind of interaction in the group? How are groups similar or different?

2. What are people talking about? How are different subjects related to each other? Who are the people interested in these topics?

66

3. Whose message to follow? How much of the group interact with each other? Who are the core people in each group? How much do they interact with each other?

This information should aid the process of selecting which group to post to and which postings to read, as well as evaluating the postings of other participants. As a result, we believe that these visualizations will encourage participants to learn more efficiently from these groups and contribute more effectively to them.

5.5.2 Basic views

Successfully completing these tasks depends on a fairly fine-gained understanding of the newsgroup(s) of interest, their topics of discussion, and the relationships between participants. These are exactly the kind of collaborative information that we have been discussing. The visualization framework we have discussed can be used to generate new views of the data to help users address these questions. Here are the views NetViz presents:

- Inter-Group View – an Object-to-Object Diagram

- Inter-Thread View – another Object-to-Object Diagram

- Thread-to-Poster View – an Object-to-Person Diagram

- Inter-Poster View – a Person-to-Person Diagram

Each view is a network diagram that can be created using a standard spring-based model for layout. Each node exerts a repelling force on other nodes, and each edge exerts stabilizing forces on the nodes it connects when its length differs from its preferred length. The preferred length is determined by the tie strength: stronger tie strength results in a shorter preferred length. The user can directly select a node, move it, and see the connected nodes and nearby nodes move accordingly to minimize the forces on them. This is an intuitive way to sort through the data and interact with the display.

In the next few sections, we will describe the basic views in more details. We will also illustrate how interactive techniques can help users utilize the basic views to answer questions about groups, threads, and people.

5.5.3 Groups

The Inter-Group View shown in Figure 5.6 displays all the newsgroups linked to a selected newsgroup (or set of newsgroups) by cross-posts, or messages sent to more than one group. This is an example of an Object-to-Object

Diagram. Two groups are closely linked if there are many shared cross-postings from people. Strong ties may indicate similarities or shared interests between two groups. This view is good for overviews of individual groups, their interconnections, and the overall context.

Say we are interested in what people think of Microsoft, and have picked two core groups of interest: alt.fan.bill-gates and alt.destroy.microsoft. The resulting diagram in Figure 5.6 contains many nodes representing various groups that have cross-posts with the core group. In other words, people who are interested in the core groups are also interested in these other groups.

Step 1. Inter-Group View

Step 2. Filtered view shows clusters of related groups

Step 3. Further filtering reveals the strong related interests of users in the alt.destroy.microsoft group.
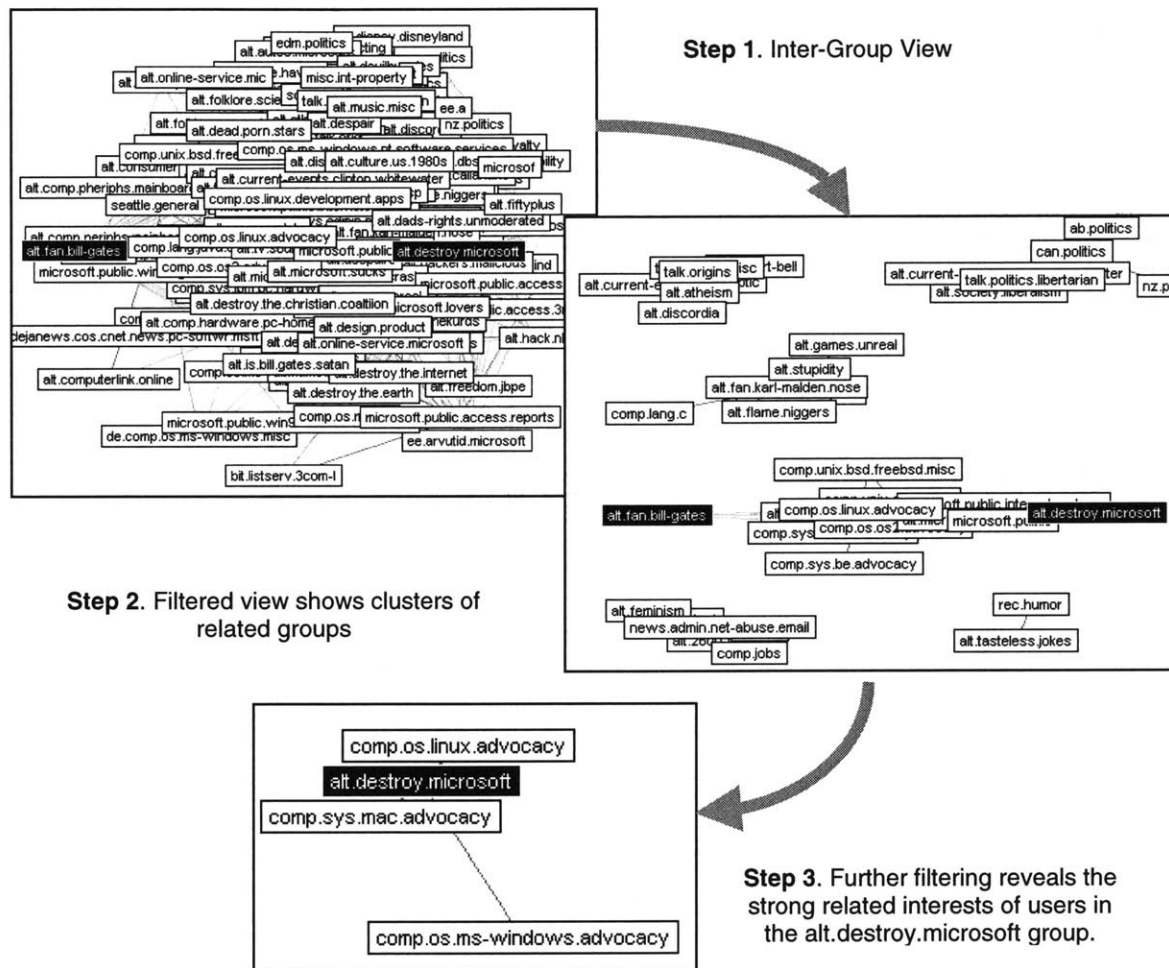
Figure 5.6 shows how Inter-group view can be filtered to show both weak and strong ties.

The first view in the top of Figure 5.6 looks rather complex. Fortunately, many of these edges are weak, consisting of only one or two cross-postings. Users can dynamically filter out the weaker links using a slider. As shown in the next view, clusters of newsgroups that tightly bind to each other start to appear. Some of these make clear sense, for example, the cluster in the upper right is about politics, and the cluster in the lower right is related to jokes, but other cluster are less obvious, suggesting interesting connections. Finally, last view in Figure 5.6 shows the strongest connections to the group alt.destroy.microsoft. It gives the user a fairly good indication of the interests of the people involved in the selected target group.

To determine which newsgroup to post or read, users need to know the amount of interaction and overall activity within each group. For that, users can overlay information on the Inter-Group View, as shown in Figure 5.7. The numbers in the parentheses represent number of poster and posts, respectively. The size of the box shows the amount of postings, and the color represents the ratio of the number of p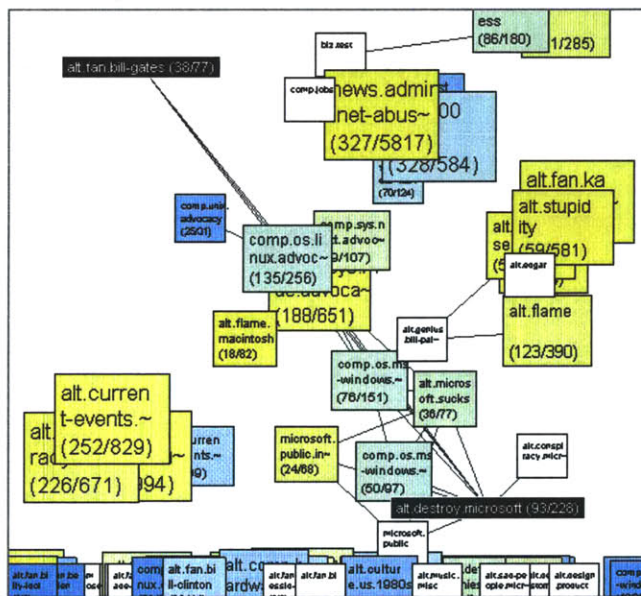osters to posts. The poster-post ratio is an important hint of the nature of intra-group interaction. A group with ratio of one (in blue) could be an announcement group since each person posts only one message -- not much interaction between people. A group with ratio close to zero (in yellow) could be a moderated group with one person or a few people posting all the messages. A group with ratio somewhere in the middle (in green) may indicate a group with healthy interaction. The blue-green-yellow color range is selected for saturated colors that still provide a good contrast to the black text.

Figure 5.7 Inter-Group View with poster/post information shows which are popular groups and how much interaction they have.

5.5.4 Threads

Once a user has pinpointed a newsgroup of interest, it is important to know what goes on within the selected group. What are the topics being discussed? Which topics are most popular? Which topics are of general interest?

The Inter-Thread View shown in Figure 5.8 displays all the threads in a particular newsgroup linked by shared posters. A thread consists of an initial post, its replies, and subsequent replies. By looking at the topics of the threads, one can get a sense of the topics being discussed within the newsgroup. This view is another example of an Object-to-Object Diagram. Two threads are closely linked if they share many of the same posters. Threads that do not share posters with any others are isolated and displayed at the bottom. These may represent topics that do not generate a large amount of interest from the newsgroup.
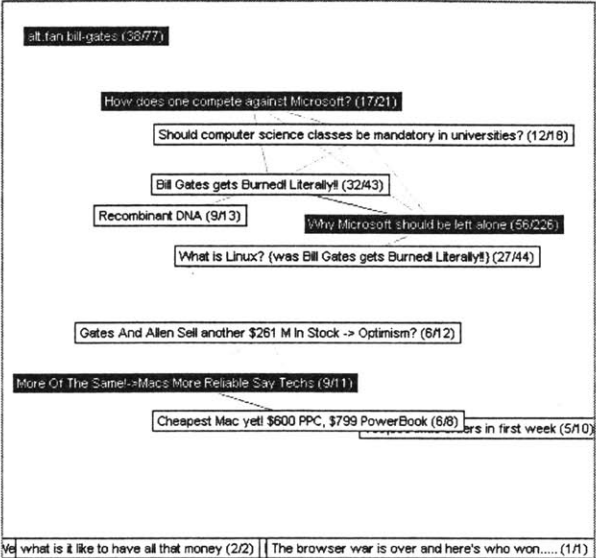

Figure 5.8 Inter-Thread View shows how topics are related.

The Inter-thread View in Figure 5.8 gives users a fairly good sense of the topics under discussion for the group alt.fan.billgates. We can see two clusters being formed: the upper one about Microsoft in general, and the lower one about Mac. By overlaying information about the number of posters and posts of the threads, we can also identify topics of hot discussion by multiple people, as in the Inter-Group View.
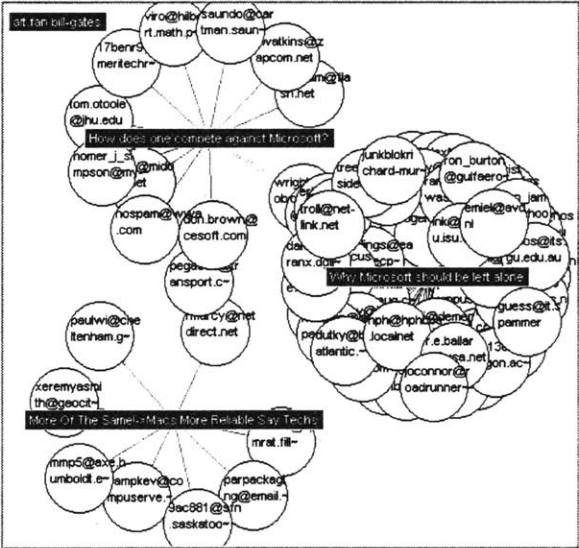
Figure 5.9 Thread-to-Poster View shows who participates in a given discussion.

5.5.5 People

From the Inter-Thread View, we can then select particular threads of interest, and look at who are the people involved in the discussion by turning to the Thread-To-Poster View in Figure 5.9. This view displays selected threads in a particular newsgroup and clusters the respective posters around each. It also shows who are the most active posters within a discussion thread, as these posters will be shown closer to the thread node. Again, we

70

can use filtering to get rid of weaker links. The Thread-to-Poster View is an Object-to-Person Diagram. It is a good way to identify people of shared interest to interact with.

We now turn to more detailed questions about people. What percentage of the group interacts with other members? Who are the core members? How often do they interact with each other? The Inter-Poster View, a Person-to-Person Diagram, is designed to answer these questions. Figure 5.10 displays selected posters in the alt.fan.bill-gates newsgroup, connected by who responds to whom.
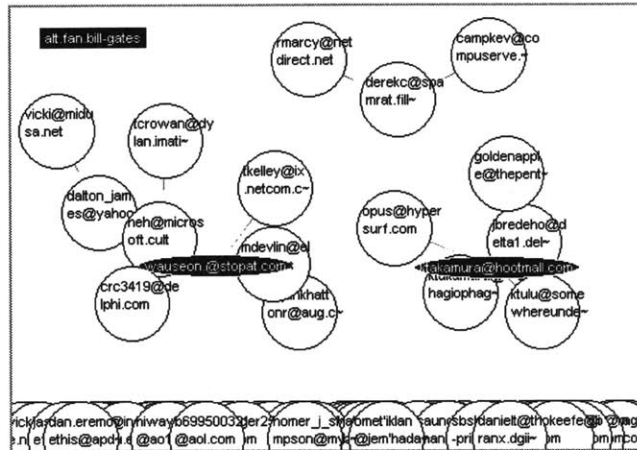

Figure 5.10 Inter-Poster View shows who responds to whom
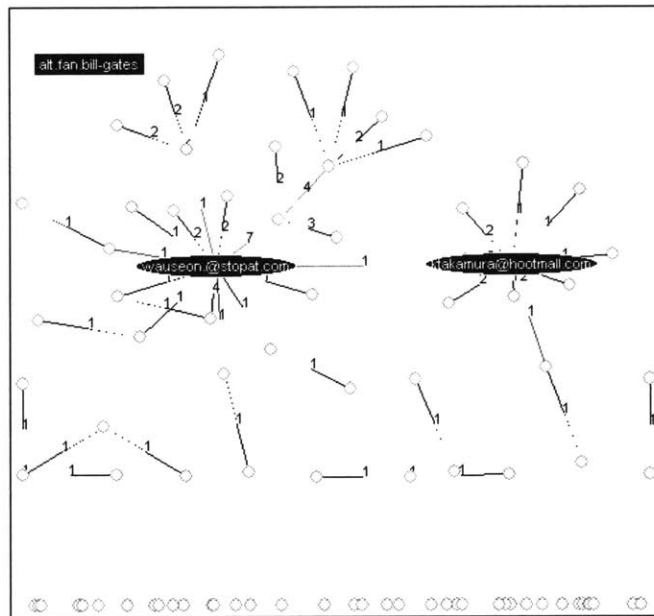
Figure 5.11 A structured Inter-Poster View shows how much interaction there are among posters.

The people who are highly connected have interaction with many people, and might thus be among the core group. Some posters post infrequently but are responded to by many others. Other people may post frequently but receive few or no responses. This view thus show the amount and type of the interaction going on in a group, and also about the status and behavior of individual participants. Figure 5.11 is another view of Inter-Poster View that shows the structure more clearly by displaying the names of only highlighted nodes, and the strength between ties.

Immediately, one can see that about half the people, actually respond to each other, represented here by connected nodes. The other half of the nodes are not connected and thus have fallen to the bottom of the display, representing

users who did not get a response. In contrast to the Inter-Group and Inter-Thread Views, which have undirected edges, edges in the Inter-Poster view are directed, showing who responds to whom. The yellow (or lighter-colored) half of the edge indicates the respondee. The visualization shows clearly that there are two hubs in the group. The one shown on the left represents a user who both responds to and is responded by others – perhaps an expert in the group; the one on the right represents a user responded to by many others – perhaps someone who started a very interesting discussion, or someone starting a flame war.
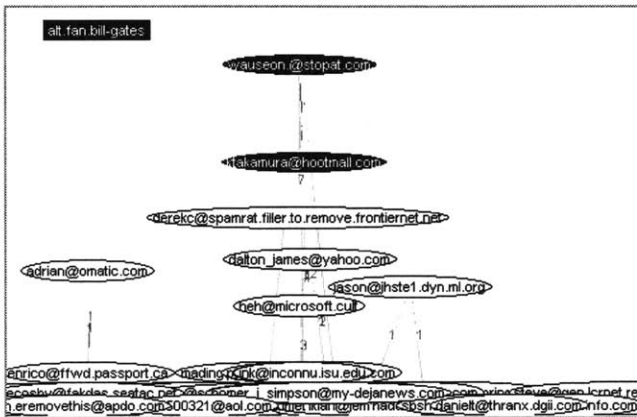

Figure 5.12 A sorted Inter-Poster View shows who are the most replied-to posters.

To clearly identify the active participants of the group, we can use the sorted display, as shown in Figure 5.12. The nodes with the highest degree, or largest number of connections, float to the top. We can immediately see who forms the core of the discussion group. Based on this information, the user can gain a better awareness of the people involved in the discussion, and become a more informed participant.

5.6 Implementation and discussion

The NetViz visualization is a Java applet publicly available from the Netscan Web site. On top of the traditional spring based model, we add for each node an ideal position. In each step of the main loop that moves the nodes by minimizing the forces on them, we also allow nodes to move at a given speed to its ideal position. By setting the ideal position to the bottom of the screen, we implement falling. By setting the ideal position to be based on degree of connection, we implement sorting.

Because our work deals with information about individual users, we must be concerned with privacy for the users. NetViz uses only publicly available information. Users who do not wish their email address to be used can protect

their identity by posting using anonymous re-mailers. Also, Netscan respects the no-archive tag in the header and does not save messages with the tag.

Our current results suggest a number of further enhancements. The layout algorithm is a simple scheme that works well for a reasonable number of nodes. When there are more nodes, a better layout scheme is needed. Graph layout is a NP-complete problem with its own research area. We hope to continue to adapt development in this area to enhance user-oriented visualization. On the more interactive side of the visualization, we would also like to allow users to navigate or go on automatically piloted "tours" of the interaction space based on a range of parameters like following the strongest, weakest or most recent tie.

We would also like to integrate more interaction, for example, pairing a synchronous chat with each Usenet newsgroup so concurrent users can interact directly while using the system. We see a way to support a chain of steps starting with searches, leading through navigation of graphical displays and culminating in direct interaction between people.

5.7 Design lessons learned

We started with the problem of how to enable users to see the connection between people and objects. We applied an existing visualization, a force-based graph, to this problem, then offered enhancements to make it much more dynamic, responsive, and useful.

As our first design experiment, NetViz has taught us many valuable lessons. First, the aesthetics are very important. While NetViz does successfully present all the data, how this data is presented is still lacking. The nodes in NetViz, boxes with text, are not very pleasing to look at. Also, the yellow-green-blue color scheme is too contrasting and unintuitive. Later versions of NetViz have improved appearance.

We also learned that motion is very engaging. People found the force-based layout inviting and enjoyed very much playing with it. A lot of people have also expressed that they like the falling metaphor for indicating nodes that are

no longer connected. Metaphors based on motion can convey information in an easily understood and visually pleasing way.

Finally, we note that for complex relationships, it might take multiple views to show the data. We see how busy displays can be filtered to show only the most relevant information. We also see how the four different views in NetViz can be used to show information at different levels and from different perspectives.

In summary, we have learned the importance of the following in visual design

- Aesthetics
- User exploration
- Motion
- Metaphor
- Multiple views

5.8 Summary

Currently, many network interaction systems are predominantly textual and require users to engage in a great deal of clicking and branching in order to assemble piece-by-piece a holistic view of the interaction space. We believe that these spaces need significant enhancement to reintroduce the social presence present in physical interactions that is missing in the on-line systems.

We have presented a framework for generating visualizations of the on-line social context. The four different diagrams (Object-to-Object, Object-to-Person, Person-to-Person, and Person-to-Object) show different interconnections between people and objects, such as books they have purchased. We introduce novel visualization techniques to reduce the complexity and increase usability of these views. As an example, we have applied the framework and interactive techniques to the Usenet groups in our NetViz system. We showed how NetViz can be used to find newsgroups and people of interest. Our framework and visualization techniques can be applied to other online social interaction environments.

6. PeopleGarden: creating data portraits for users

In chapter 3, we discussed that understanding the patterns of behavior a given environment is an essential step for successful interaction. How do we go about gaining such an understanding? Many on-line interaction environments have a large number of users. It is difficult for the participants, especially new ones, to form a clear mental image about those with whom they are interacting. How can we compactly convey information about these participants to each other? We propose here the *data portrait*, a novel graphical representation of users based on their past interactions. Data portraits can inform users about each other and the overall social environment.

6.1 Introduction

On-line interaction environments such as mailing lists, Web-based message boards, chat rooms, and Usenet newsgroups have become widely popular. How can we figure out who is who in these environments? Gestures or facial expressions that convey rich meanings in real life interaction are missing on-line. At the same time, a much larger number of participants can directly interact with each other on-line, unlike in real environments. It becomes increasingly difficult to distinguish individual users and to comprehend the overall interaction context. Often, users must piece together a holistic view of other users through extended participation. This makes it difficult for new users to find appropriate groups to participate in and people to interact with. Even existing users may find it difficult to keep track of an ever-changing population.

What is missing is effective communication of information about participants to each other. Figure 6.1 lists four questions users might ask about the underlying social environment (the quantitative measures in parentheses provide insights to these questions):

1. Do participants here really get involved? (Post frequently or only once in a while?)
2. Who are the core members? (Who has been here for some time or posted many messages?)
3. How much interaction is there? (Do people respond to each other?)
4. Do participants here welcome newcomers? (Do newcomers get many replies?)

Figure 6.1 Users' questions about a discussion group.

To answer these questions, one needs to be able to examine multiple pieces of information at once, such as how much one posts, how much response one receives, or when were the messages posted. A naive approach to gain information about users would be to sort all the postings by author [Deja99]. While such views offer some information about a single author's level of involvement or interests, users still need to search for responses from other authors or for comparison with other authors of the group. We need the ability to look at different posting and response patterns for all the members of a group together. Such patterns are hard to see by viewing lists of values or tables of data. However, by visualizing the salient data, the key roles and relationships can be readily perceived.

Our solution is the *data portraits* (see Figure 6.2), which are abstract representations of users' interaction history. By allowing participants to see the past behaviors of other users, we can help them gain an idea of how to appropriately interact with them in the present. These visualizations satisfy two important goals:

1) Providing rich details to illustrate individuals

2) Being compact enough to allow an overview and comparison of a group of individuals.

We will divide our discussion in this chapter into two parts. First, we discuss means for visualizing individuals, present the notion of data portrait, and discuss how it can be created using a flower metaphor. Then we discuss issues in visualizing a group of people, and present our garden metaphor. We will see how data portraits like the ones in Figure 6.2 can be used to illustrate an interaction environment.
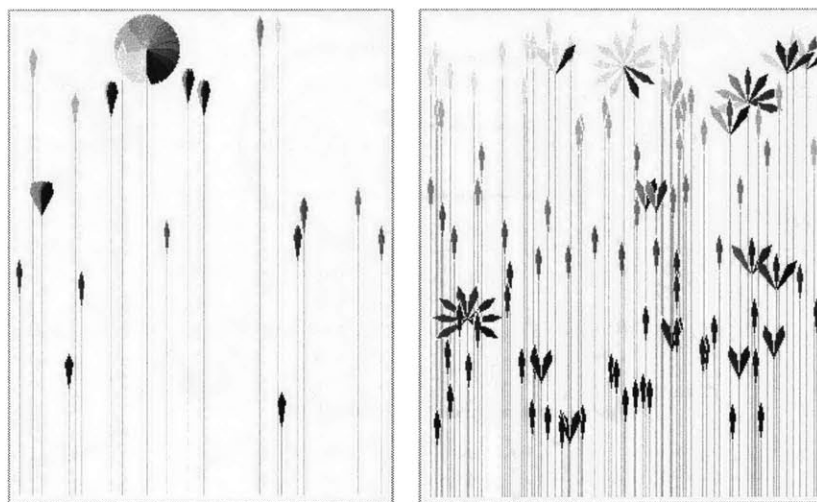

Figure 6.2 What can data portraits like the ones above tell us about discussion groups?

6.2 Visualizing Individuals

Visualizing people is not new: people have been painting portraits for ages. Successful portraits not only show the physical likeness of their subjects, they also use objects in people's lives to illustrate their interests and status [Brilliant93]. Figure 6.3 shows three works by various well-known painters: (a) Vincent van Gogh's self portrait of an artist at work, (b) a famous work by John Singer Sargent, one of the finest portrait painters of our time, depicting children in an opulent but imposing surrounding, and (c) a Pierre-Auguste Renoir's painting of a mother and child in warm and peaceful surroundings.
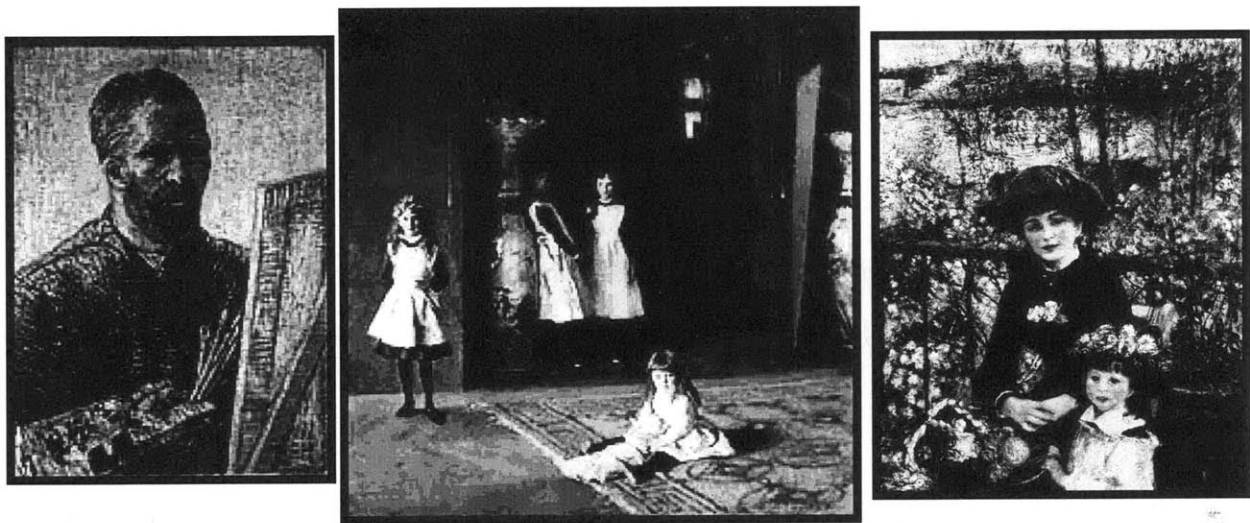
(a) van Gogh, 1888 (artmuseum.net)   (b) Sargent, 1882 (www.boston.org)   (c) Renoir, 1881 (artmuseum.net)
Figure 6.3 Portraits by various famous artists illustrate the use of objects to present the subjects' life.

Algorithmically generating such telling portraits of people may be impossible with current technology. However, we can borrow a useful insight: we can use objects from peoples' on-line life to illustrate their positions and status.

6.2.1 Avatars and other graphical representations

More recently, avatars have been used to represent individuals in virtual environments. Avatars include such varied representations as smiley faces, cartoon characters, 3D models, or photographic representation of users [Damer97, Kurlander96]. These avatars have been used in a variety of 2D and 3D chat rooms and virtual environments.

Most graphical chat rooms allow users to select a pictorial representation of themselves. More interesting are avatars that can show a limited range of information about the users' current state. Steve Benford's 3D avatars show where in the virtual environment a user is looking at and acting upon [Benford94]. Judith Donath's Illustrated Conversation shows by gaze direction who is talking to whom in a conversation [Donath95b]. Fernanda Viegas uses a more abstract visualization. Her Chat Circle represents users as circles, and changes the size of a circle as the user sends out a message
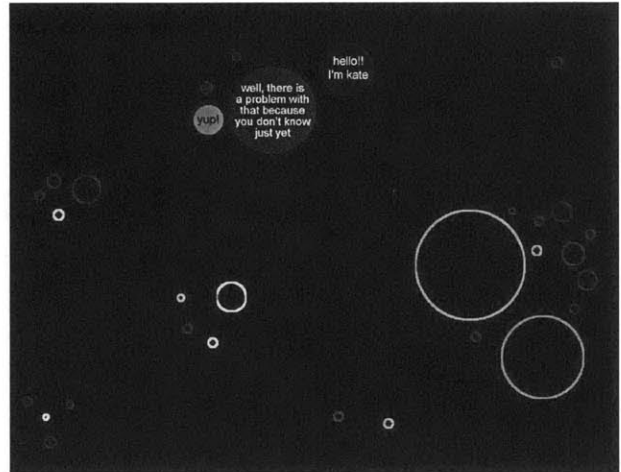

Figure 6.4 ChatCircles represents individuals abstractly as circles, and enables the formation of conversational groups [Viegas99].

[Viegas99]. Users can move their circles to form conversational groups, as shown in Figure 6.4. Vertegaal et al's GAZE system (Figure 6.5) uses videos of the individuals to represent them, and also shows their different focus in the shared document [Vertegaal99].
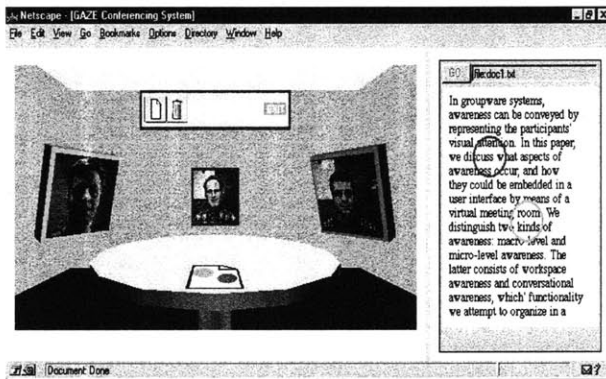
Figure 6.5 GAZE system shows video avatars of individual users and indicates where in the shared document they are currently viewing [Vertegaal99].

These avatars can show a range of current states of the individuals, but do not convey much about the past behavior of users. They do not show how long the person represented has been in an environment, or whether he/she is a prolific contributor to the conversation. More information about the users' overall pattern of behavior is needed.

There has been relatively little work on visualizing historical information about individuals. One notable exception is Plaisant et al's LifeLines, which shows past history of a person using time-lines [Plaisant96]. It is effective for detailed examination of an individual's medical or criminal records, as shown in Figure 6.6. However, it is not designed for a quick summary of a person or for comparing a group of people.

6.3 Data Portrait

Figure 6.6 LifeLines uses timelines to show detailed personal history of a person [Plaisant96].

While there are previous visualizations that convey a user's history or are compact enough to represent individuals in an on-line environment, we have not found a single visualization that combine both. Our solutio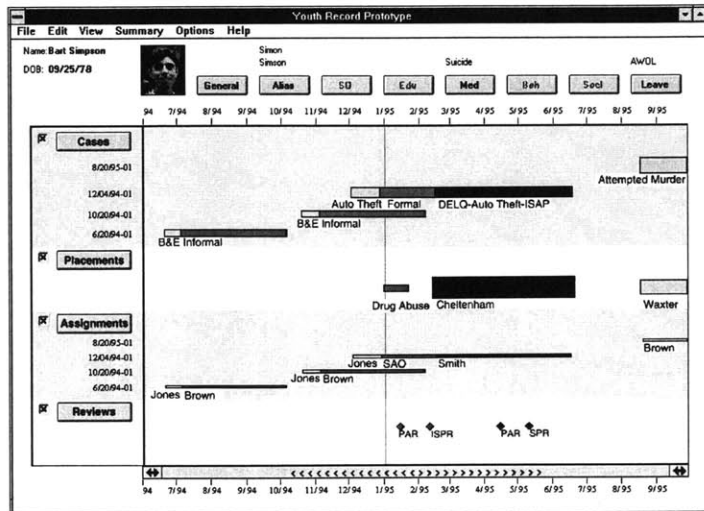n is the *data portrait*, or user visualization based on interaction data, both to represent individual users and to compare multiple users.

We are inspired by how painters use objects in people's lives to illustrate their interests and status. The on-line equivalents of one's objects are data about one's past interactions. For a Web-based message board, where a user can post messages or respond to others' messages, his data objects are the set of messages posted by him. In a chat room, the data objects are the chat messages sent by each chatter.

A data portrait is a visualization of a user based on his data objects. Different portraits can be created for different users depending on the number and attributes of the data objects. Possible attributes for data objects like messages include time of creation, size, addressee, whether an initial message or a response, or the number of responses from others. Various visual parameters such as the size, hue, saturation, or shape of the data portrait can then be used to represent these attributes. The challenge in creating successful data portraits is thus two-fold: 1) selecting those attributes about data objects that best convey information about a user, and 2) deciding how to visualize these attributes legibly and intuitively.

Over time, a user's portrait can change based on his interaction. The portrait will thus always show the most up-to-date data about the user. Unlike avatar representation, these portraits are based on real interaction history. A collection of data portraits can be used to represent the participants of an on-line environment. Such an overview allows users to see overall trends in the data, such as the level of user participation or the amount of interaction.

79

Because data portraits are generated by simple, uniform rules, they form small multiples that can be compared and contrasted [Tufte91].

6.4. PeopleFlower

We will now discuss in more detail PeopleFlower, our visualization of the individual. Each user is represented as a flower, as shown in Figure 6.7. His data objects are represented as petals of the flower. As more data objects are created, they are added clockwise to the flower. The entire flower is then re-centered to maintain symmetry.
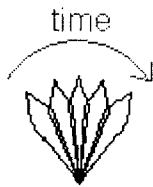
Figure 6.7 PeopleFlower.

Why use the flower metaphor? Visual artists have long emphasized using representations that have the same level of visual and conceptual complexity as the objects they represent [Arnheim89]. We therefore wanted a simple object that can easily deal with a changing number of components. We also like the organic nature of a flower, and the suggestion that it changes over time, as users do. The flower metaphor was chosen for its simplicity and intuitiveness. Other visualizations that use radial layout to display data include radar sweep visualizations and Kiviat Diagrams [Jefferies97, Harris97]. For more technical discussions on creating organic or non-photo-realistic rendering, see [Kowalski99, Raskar99]. For simplicity, we will focus our discussion on a Web-based message board. Our findings can be generalized to visualizations of Usenet newsgroups, Web page traversals, or Chat room interactions with few modifications, as discussed in Section 6.6.

For a message board, each PeopleFlower represents a user who posts at the message board, and its petals his postings. The different visual parameters of a petal can be used to represent different attributes about each posting. We found the time of posting, the amount of response, and whether a post starts a new conversation to be three valuable attributes in conveying social information about the user, as demonstrated below.

6.4.1 Time since posting

Time of posting is an important attribute for showing posting history. PeopleFlower indicates time in both the ordering and the saturation of the petals. A message posted far in the past would be less saturated and less noticeable. As a user posts more messages to the board, his or her PeopleFlower will change, as shown in Figure 6.8.

Figure 6.8 A user's PeopleFlower changes over time.

A few things to note here:

- The number of petals increases as more messages are posted, following the metaphor of a flower booming.

- Older petals move to the left as newer petals are added to the right, to maintain symmetry. The overall shape of a flower is determined by how many petals it has.

- Each petal fades over time showing time since posting.

- A marked difference in saturation of adjacent petals, seen in the rightmost flower, denotes a gap in posting.

6.4.2 Response to posting

PeopleFlower can display not only a user's own posting pattern, but also the amount of feedback from other users of the board. This is important for illustrating the amount of interaction in an environment. We have used pistil-like circles on top of the petals to denote responses. Each circle indicates one response to a post, as shown in Figure 6.9.
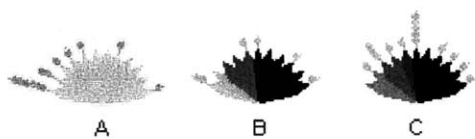
Figure 6.9 Three users with same amount of postings over different durations, and with different amount of responses.

While the three users in Figure 6.9 posted the same amount of messages, they posted at different time interval, as indicated by the saturation of the petals. The amount of responses also varies greatly.

- User A posted a series of messages in a short interval, and has not posted for a while.

- Initial response to A was high, over time it decreased.

- Users B and C both posted their messages over longer intervals, but C has gotten much more response.

6.4.3 Initial post vs. reply

Why do the amounts of responses differ for these three users? Figure 6.10 shows the same PeopleFlowers with petals colored by whether the post starts a new conversation (i.e. initial posts versus replies). This scheme has been used in the Loom visualization for newsgroups [Donath99].
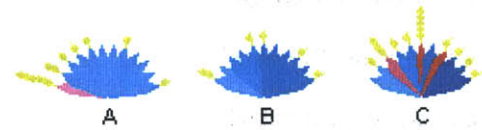
Figure 6.10 The same three users from Figure 5, this time with magenta representing initial posts, blue replies.

We gain more insight into the situation:

- User A started with an initial post (shown in magenta), perhaps a question or a new topic of discussion. There were many responses. Then User A replied to others' postings, and the level of response gradually dropped.

- User B only responded to others' posts, as indicated by all blue posts.

- Users C made both initial posts and replies. This user likes to both start new conversation and interact with others. User C received the most replies.

This example seems to suggest that one way to get more response from others would be to start a new conversation. See [Whittaker98, Smith99] for more studies about the interaction dynamics of Usenet newsgroups.

6.5 Visualizing relationships among people

Having dealt with the visualization of individuals, we now look at how to visualize a group of people. Social network diagrams have long been used to illustrate the relationships among people. In these diagrams, as shown in Figure 6.11, each node represents a node, links between nodes represent various connections between people [Moreno36, Wellman79, Krebs96, Kautz98]. These diagrams can get fairly complex as the number of nodes rises, and are better suited for sociological studies rather than quick overviews. Our own NetViz makes use

Figure 6.11 Social network diagrams illustrate the relationship among people, such as co-authors of a paper [Kautz98].

of such diagrams, as described in the previous chapter, and simplifies it for the average users.
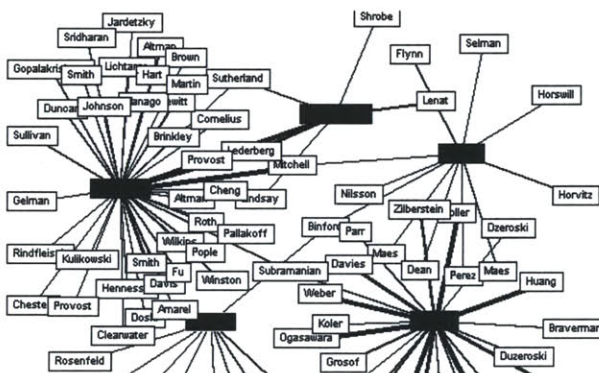
A related piece of work is Judith Donath's Visual Who [Donath95], shown in Figure 6.12. It shows the people's affiliations in various communities. A set of mailing lists are selected as anchors, the names of people are then placed around the anchors based on how similar they are to members of these lists. This removes the messy links between nodes and still shows their inter-relationships.
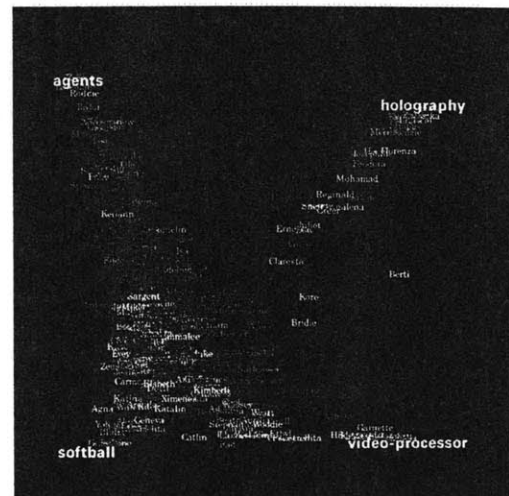
Figure 6.12 VisualWho illustrates the group affinity of [Donath95].

While these visualizations can convey quite a bit of information about relationships among people, they do not show much detail about the individuals. It is still difficult to gain a comprehensive view about these individuals.

6.6 PeopleGarden

Our PeopleGarden visualization addresses this problem by showing the interaction history of the individuals. We make use of the PeopleFlowers described earlier.

Figure 6.13 shows a section of a sample board with about 1200 messages posted by 150 authors over the period of two months. Currently, users need to read through a large amount of text to determine who's who in the message board. PeopleGarden is designed to facilitate this task. While PeopleFlower specifies

```
2 methods - Cindy 17:41:20 4/24/99 (6)
    Re: 2 methods - Kai 20:55:38 4/26/99 (2)
        Re: 2 methods - Henry 09:54:09 4/28/99 (1)
            Re: 2 methods - David 19:02:58 4/29/99 (0)
        Re: 2 methods - Pierece 14:42:10 4/26/99 (0)
    Re: 2 methods - David 10:01:51 4/25/99 (1)
        Re: 2 methods - Dirk 19:58:21 4/25/99 (0)
Apology for "provoking" - Cathy 22:00:48 4/22/99 (2)
    Re: Apology for "provoking" - James 12:47:31 4/23/99 (1)
        Re: Apology for "provoking" - Mia 06:05:28 4/24/99 (0)
L/R Brain info. continued... - Kam 18:44:28 4/22/99 (1)
    Re: L/R Brain info. cont... and.... - Cathy 21:02:30 4/22/99 (0)
```

Figure 6.13 Traditional view of a message board.

how to represent individual users, PeopleGarden specifies how to layout the flowers for overview and comparison of users.

We have used the garden metaphor because a healthy garden has certain properties that we can use to represent a healthy discussion group. For example, a garden with more bright flowers indicates a discussion group with more new posts. It looks healthier than one with faded flowers representing a group with mostly old posts. Figure 6.14 shows a PeopleGarden created for the same message board. The height of a flower represents how long a user has been in the board, as indicated by time of first posting. This uses the metaphor that a flower planted earlier is taller than one planted later. A few patterns here:

Figure 6.14 A PeopleGarden showing messages from a message board with 1200 postings over a 2-month period. Height of flower denotes amount of time a user has been at the board.

- Many people have participated at the board since the beginning, and have been prolific contributors since.

- A large number of casual participants dropped by to make a few posts.

- There are more users joining the board in the first half than the second.

We can also notice outliers such as the big red flower two-thirds from the top and two-thirds from the left of the figure.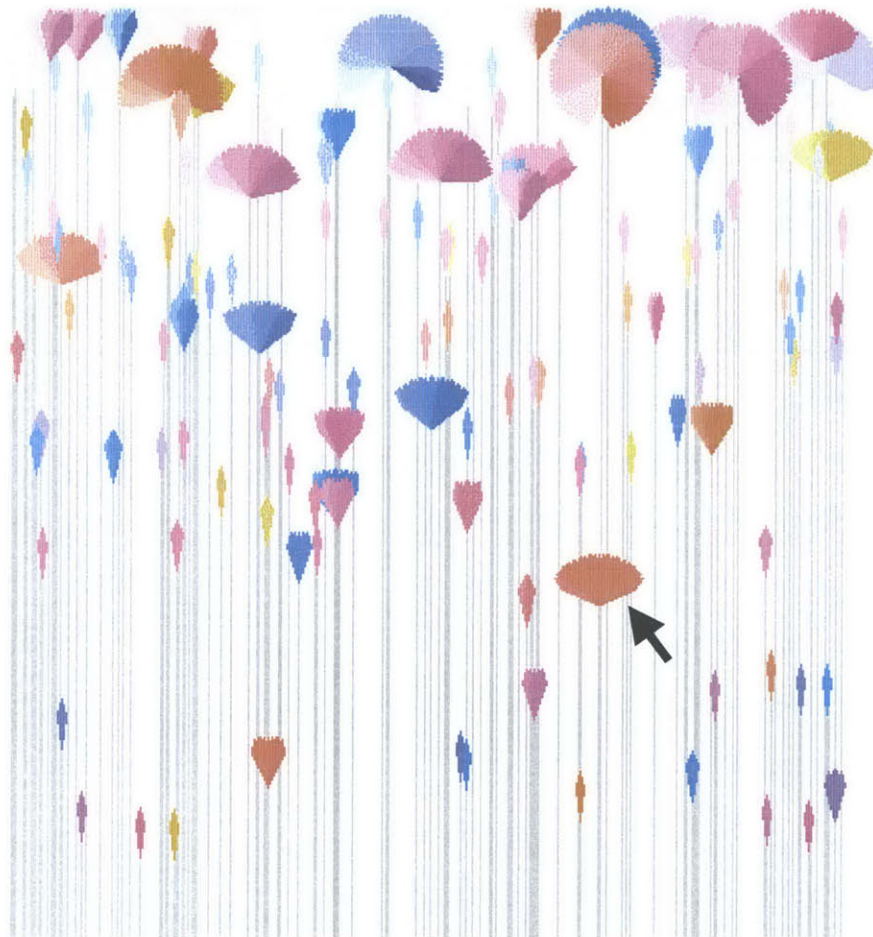 While most people have more to say the longer they have been at the board, this relative latecomer has really bloomed and contributed much to the discussion.

Now let us return to the four user questions we have raised at the beginning of the paper about on-line environments:

Q1. Do participants here really get involved? (Post frequently or only once in a while?)

Q2. Do participants here welcome newcomers? (Do newcomers get many replies?)

Q3. How much interaction is there? (Do people reply?)

Q4. Who are the core members? (Who has been here for some time or posted many messages?)

The PeopleGarden in Figure 6.14 can help to answer a couple of these questions.

A1. There are many users here with varying level of involvement. Many of them have also posted continuously, denoted by smooth change of saturation across a wide range.

A2. The most prolific contributors also tend to be those who have been here the longest, and can be found at the top of the figure.

Figure 6.14 is useful for answering questions about general posting trends, since all the flowers are displayed at once. For detailed questions about the level of responses to individuals, another view is required.

Figure 6.15 shows a different PeopleGarden that lays out flowers sorted by the amount of petals. Posts are colored differently based on whether it is an initial
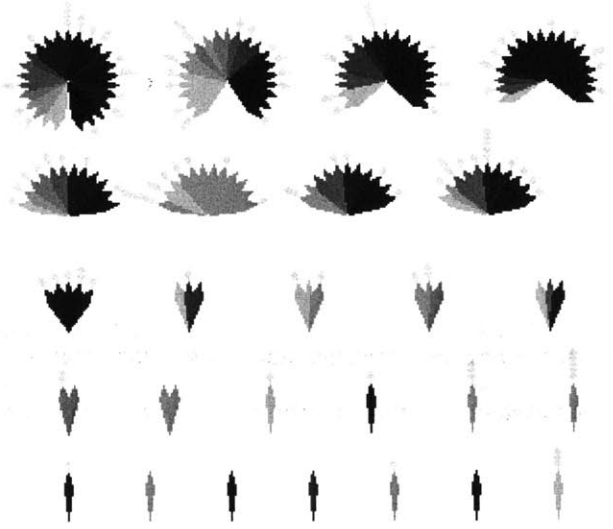
Figure 6.15 A PeopleGarden sorted by amount of postings. Magenta denotes initial postings, and blue replies.

post. The responses for each post are also shown. A smaller message board is used here to show the details more closely. This figure can help to answer the remaining questions:

A3. People in this board respond to each other quite a bit. In fact, most posts are responses shown in blue, and more than half the posts are replied to, as indicated by the gold pistils.

A4. Newcomers shown at the bottom have a 50-50 chance of getting a response, only slightly worse than the old-timers. Also, their initial posts are much more likely to get a response.

6.5.1 Comparing different groups

PeopleGarden applied to multiple groups shows how the interaction in these groups differs. For example, Figure 6.16 shows the PeopleGardens for two groups with around 200 posts each, but very different interaction. Height of a flower again denotes how long a user has been at the board.
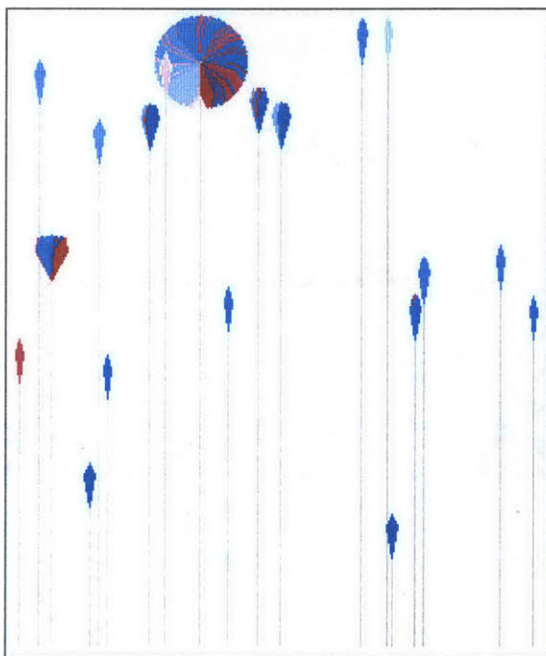
Figure 6.16 (a) A group with one dominating voice (http://graphics.lcs.mit.edu/~becca/enneagram/movieboard,July '99) with a large number of replies in blue. Height of a flower denotes how long a user has been at the board.
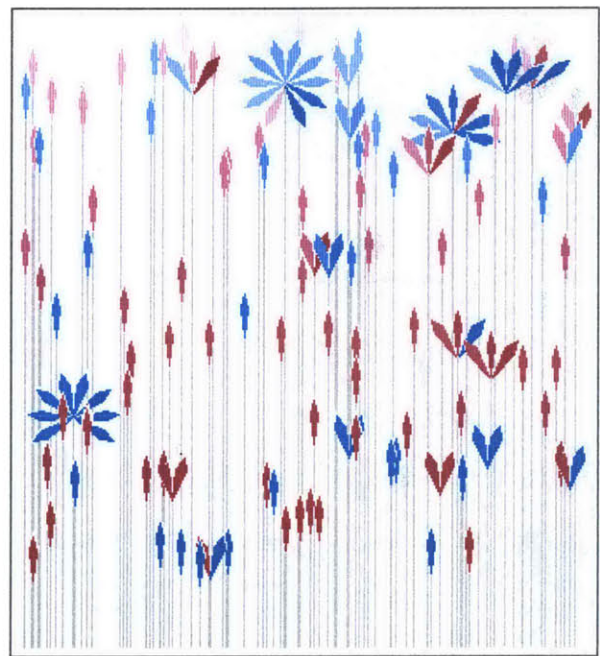
Figure 6.16(b) A more democratic group (http://www.usscreen.com/message/general/, July '99) with a large number of initial posts in magenta. Height of a flower denotes how long a user has been at the board.

Figure 6.16(a) shows a group with one dominating voice with many initial posts (in magenta) and more responses (in blue). There is a trickle of new participants throughout the time period, most of whom tend to post replies (in

blue). There is a trickle of new participants throughout the time period, most of whom tend to post replies. This is more a discussion group with predominantly replies in blue. Figure 6.16(b) shows a more democratic group with more equal participation. There are a larger number of participants, each of who posted between 1 and 10 messages. People tend to post more initial posts than replies. This suggests that this group is more of an announcement group.

6.6 Discussion

PeopleGarden has been implemented as a Java applet. Each petal of a PeopleFlower is a 4-sided polygon. To create a flower, draw each petal in a clockwise fashion. The flower with the largest number of petals almost completes a circle. It determines the spacing between petals for all the flowers in a garden. PeopleGarden lays out the individual flowers by setting their center point. For example, the sorted view reduces row by row the spacing between the center points of flowers. Thus, the later flowers with less information will not take up disproportional amount of screen real estate.

How does PeopleFlower scale as the number of posts for a given user increases? We found that the petals can still be distinguished clearly when the ends of adjacent petals are more than one pixel apart. However, pistils representing responses cannot overlap, or the result would be illegible. One solution would be to use each petal to represent the average or median value of multiple data objects.

The PeopleGarden visualization scheme can readily be applied to other threaded discussion space such as Usenet newsgroups. It can also be applied to chat rooms with some modifications. One small difference is that there is no clear distinction between an initial post and a response. We can use instead different categorizations of chat messages such as a question versus a statement. A related application is representing users' traversal history through a Web site, as in the LiveWeb system described in the next chapter. We would like to dynamically update PeopleFlower visualization of users. One can watch the growth of a PeopleGarden over time as users interact more.

Since our visualization is based on users' actual history, we need to be concerned with their privacy. The visualizations presented here use only publicly available posting histories. For more sensitive information, a possible

solution is to not specify user identity, or use only pseudonyms. Our visualization can still show patterns about an interaction environment without identifying the individuals.

6.7 Design Lessons Learned

While we have not conducted formal user testing, we have collected informal feedback on PeopleGarden. People responded enthusiastically to its visually pleasing and organic qualities. They also found the visual encoding of information to be fairly easy to understand. People prefer encodings with some basis in reality. For example, most users found it easy to tell how long a user has participated in the message board based on the height of the flower. For encodings not based on physical reality, visual simplicity is essential [Tufte90]. For example, the two-tone coloring of petals is fairly easy to comprehend. But when more than two colors are used, the display becomes too cluttered. Also, we have experimented with showing time of posting based on angular position of petals, but this produced irregular flowers that were too complex to compare at a glance.

We received feedback about from other interface designers at the User Interface Software and Technology Conference where PeopleGarden was presented [Xiong99]. Using an organic metaphor like flowers to represent people is an effective approach that leaves rich room for other related metaphors. One researcher suggested for example that bees flying around could be used to illustrate lurkers, who read the postings but do not post themselves.

Some researchers suggested the use of other visual representation schemes for creating data portraits. We agree very much that the flower garden is only one of many possible ways of representing people. Other schemes include a school of fish, a colony of ants, or a flock of birds. We prefer organic schemes for the rich set of metaphors they can offer: the three alternative schemes we proposed are all based on animals that tend to move together as a group, but each also has its own characteristics. It is also important to balance between visual appeal and the ability to convey information. One should not get lost in the details of rendering a beautiful bird if it conveys little additional information to the user, or worse, if it reduces legibility. The flower scheme we have used, for example, offers a visually pleasing representation that is at the same time quite legible.

In summary, we note the importance of

- Aesthetics
- Natural metaphors
- Visual simplicity

6.8 Summary

To interact effectively on-line, one needs to know about other participants. Unfortunately, existing interaction environments require extensive participation for users to form a holistic view of the interaction context. Re-introducing the social context present in physical interactions can enhance these environments.

Data portrait is a novel method for conveying information about participants of an on-line environment. PeopleFlower visualizes a user's posting history: PeopleGarden uses a collection of these flowers to represent users of a message board. The various views of PeopleGarden can be used to successfully answer user questions about the board. PeopleGarden can be extended to other on-line environments to illustrate their social context. These data portraits can help users to understand an interaction space, see how it changes over time, and provide a mirror for the level and nature of the interaction in that space.

Color pictures and an actual demo of PeopleGarden can be found at

http://graphics.lcs.mit.edu/~becca /pgarden

Part III: System and Experiment

In the previous sections, we have discussed the problem of lack of social presence on-line, and described several experimental designs addressing the problem. In this section, we will put our ideas to the test.

Chapter 7 describes LiveWeb, a real-time Web-based system that incorporates elements from the three designs presented in the last section. It is a more practical testing ground for our ideas about visualizing information structure and social presence for a ramp-up of interaction. LiveWeb is an experimental system to

- Collect user presence information for visualization

- Visualize the site structure

- Visualize people currently visiting the site

- Enable users to browse while interacting with other users

How can this work be evaluated? In our introduction, we proposed showing both that the problem is real and that our solution is sufficient by conducting an experiment using LiveWeb. Chapter 8 describes the experiments, which show that users indeed strongly prefer the display of presence we have created.

7. LiveWeb

How can we allow users to become aware of each other's presence on the Web? How can we enable them to see crowds of other users, watch what others do, and ask questions of fellow users? Our solution is LiveWeb, which displays a hand-created map of a Web site and then overlays real-time user traversal. Visualizing the Web site structure is necessary for providing a context in which to display activities of other users, but more importantly can help users understand and interpret others' activities. A visual representation of users allows them to become aware of others moving within the site. This flow of activities may attract other users and give rise to spontaneous interaction, just as in real life. We enable users to communicate with other users who may be involved in similar activities. We have thus broken down the problem into three parts:

1. Visualizing the structure of the site

2. Visualizing the presence of other users within the site

3. Allowing users to interact with other users

Eric Brittain, a fellow PhD student in the MIT Computer Graphics Group, has worked with me on the system engineering part of LiveWeb, and also contributed ideas to its visual design.

7.1 Our approach

We believe that displaying presence on the Web is important because its information content can serve as a meaningful context for social activity. What connects two Web users from different locations and backgrounds may be a Web site on a shared interest, or a page addressing a shared concern. In fact, the more narrowly focused the content of a Web page, the more meaningful it becomes when two users are viewing the same page. Making the informational structure visible is thus crucial to showing user presence, facilitating understanding of others' interests, and connecting people on-line.

LiveWeb first displays a navigable map, which can help users understand the informational structure of the site. LiveWeb makes use of hand-created maps because it provides a tailored view for the individual site, and is fairly easy to implement, allowing us to focus on the presence display. Many sites already have hand-created maps to help

users navigate. They typically list the first two to three levels of the site, using color, placement, or font size to indicate the different levels or sections of the site [Kahn99]. We will see examples in the next section of how we create maps that are appropriate for different sites.

On top of a site map, we overlay real-time user presence. Each user is represented by an icon. As they move through the site, their icons move correspondingly through the map. Thus, one can see crowds forming at various parts of the site or busy traffic at the site.

We believe this visualization serves several functions. It allows users to

1. Navigate the Web site from the overview map.

2. Identify pages that are popular.

3. See users move in real time within the site.

4. Gain a sense of the activity of other users of the site.

5. Identify potential interaction partners.

6. See a "Community Mirror" [Borovoy98].

1. Navigation: We display an overview map of the site that reminds users of different parts of the site and allows users to navigate the site using the map. Today's site maps typically list all the first or second level sites, and are often designed for a one-time overview or for finding a particular part. The function we are interested in is currently served by the navigation bar, which can be found on almost all major sites today, but lists only the main areas or functions of the site. Our maps need to balance between the completeness of existing site maps and the succinctness of navigation bars.

2. Hot spot identification: By displaying the presence of visitors on the overview map, our visualization can be used to identify which pages get a large amount of traffic. These indicate pages of general interest to other visitors of the site, and therefore could help a new visitor determine where to go, or an existing visitor determine what else to see. This function has not been explored by existing tools, which tend to focus on showing who else is on the same page.

Our visualization also has an option to show accumulated traffic in addition to current traffic. Depending on how dynamic the content is, it can be useful to show what is hot this moment or over a time period.

3. Real time movement: The site map provides the basis for displaying user movements within a Web site. As a Web visitor moves from one page to another, his icon will move accordingly. The result is a dynamic view of activities taking place at the site. By showing each individuals' movement, we allow visitors to see their movement in the context of the site overall, and of other users also surfing the site.

4. Knowledge of other users' activities: People-watching in real life is fun because we see people engaged in various businesses. Similarly, on-line crowds should show different people going about their own activities, providing a rich on-line "sidewalk ballet" [Jacobs61]. While we do not have personal information about users (and may not want to show it, for privacy reasons) we can show information about their activities at the site. We can give people different visual representations based on how many pages they have seen, or how long since they have moved.

5. Identification of interaction partners: The visualization can serve as a basis for finding others to interact with. Some users may have generic questions that they would like to ask of anyone present on the site. Others might be more selective in finding people on the same page or a similar path. Our visualization allows users to see at a glance who is on the same page, and to find out which path they have taken.

6. Community mirror: The visualization can help a visitor to gain a sense of the community of visitors at the site, of which he is a member. The visitor can view the number of users and their different behaviors to reflect on the site, or how the site and its visitors change over time: why some areas have become less popular, why there are sudden change in traffic to the site.

7.2 Visual Interface

The interface that the users see is a two-frame browser window shown in Figure 7.1. Users can view the visual site map with presence and chat in the left frame and browse the Web as usual in the right frame. They can use the visual map to select what to browse and read the details in the other frame. As a visitor browses, new people coming to the

Web site may attract his attention and pull him out of the browsing mode to chat with others. As he interacts with others, he can see what others are reading and follow their path, or see what the popular destinations are and go there. Users can therefore move freely between these two views and these two modes of interaction.
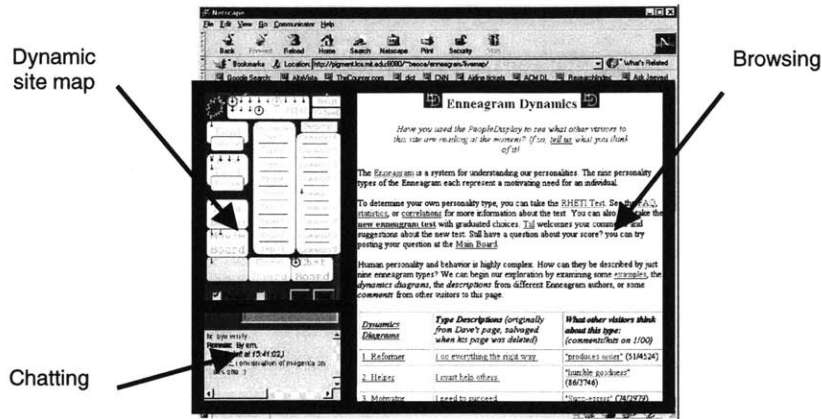
Figure 7.1 LiveWeb's two-frame view: A user can view a dynamic site map with presence and chat in the left frame, and browse actual pages in the right frame.

The two-frame model was chosen for its simplicity in achieving both the overall and detailed views for a user on an existing Web browser. However, our model does place certain design constraints. For example, the site map must be compact enough to fit in the left frame and not take up too much screen space, yet be detailed enough to be useful as an overview and a navigation aid.

7.2.1 Map

LiveWeb offers a hand-created 2D site map (as in Figure 7.2) as the basic visualization. It is useful for mapping sites with a small number of clearly distinguishable sections and shallow link structure, or for gaining an overview of larger sites. Different regions of the map represent different parts of a Web site, which can be 1) a page 2) a set of related pages, or 3) a directory of pages. Web site administrators can create a map for their own site by specifying a background picture and a mapping between a region of the map and the Web pages it represents.

Here are some simple heuristics for creating site maps:

- Determine the relevant areas of the site

- Distinguish different regions of the site using placement/appearance

94

- Use the same style for pages with similar functions

- Cluster related pages

- Emphasize important parts of the site (more popular, more content, etc..) by greater areas, distinguished appearance

Additional heuristics allow the creation of maps that can better exhibit the activity of people:

- Show popular paths through the site

- Show pages that can better distinguish the interests of visitors

We will use an example to explain these rules in more detail.
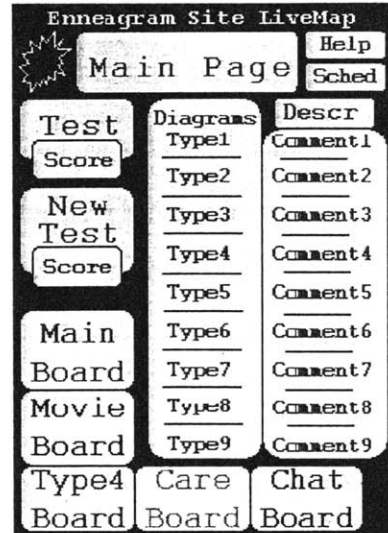

Figure 7.2 A Map of a content Web site created by hand.

Personality Site -- This site offers a couple of personality tests, diagrams and user comments about different personality types, and a set of message boards where users can post their thoughts on various subjects. The map for this site, shown in Figure 7.2, reflects the different areas. Below the main page area, we see on the left the two tests. They each have two parts: the test itself and the scoring. Their similar visual representations indicate their similarity. In the middle is the region for the diagrams, and to its right, the comments. Interest in the nine different types is an important indicator of difference among visitors to this site. The nine diagrams and the nine comments are therefore aligned to better showcase people who move between them to find out more about a particular type. Finally, towards the lower left of the map are five boxes for the five message boards. Many repeat visitors spend a large amount of time checking out these message boards to post new thoughts or reply to others. These boards are colored differently, to better indicate their difference; corresponding boxes on the map are colored to match the boards.

7.2.2 People

Our visualization overlays the map with icons that represent individual visitors to the Web site, as shown in Figure 7.3. Visitors to the site can identify which areas of the site are busier by noting where these icons tend to gather. To provide visitors with some knowledge of individual users and hence facilitate identifying potential interaction partners, we offer an indication of:

- How long a user has been on the site

- How long since he last moved

- How many places he has been

- Which places and in what order

This information can be represented using different metaphors. We chose a modified version of the PeopleFlower metaphor from the previous chapter, because it allows us to convey a lot of information compactly. (See the previous chapter for more discussion on ways of representing people.) When the visitor first enters the site, he is shown as a flower with a single petal. A new petal is adde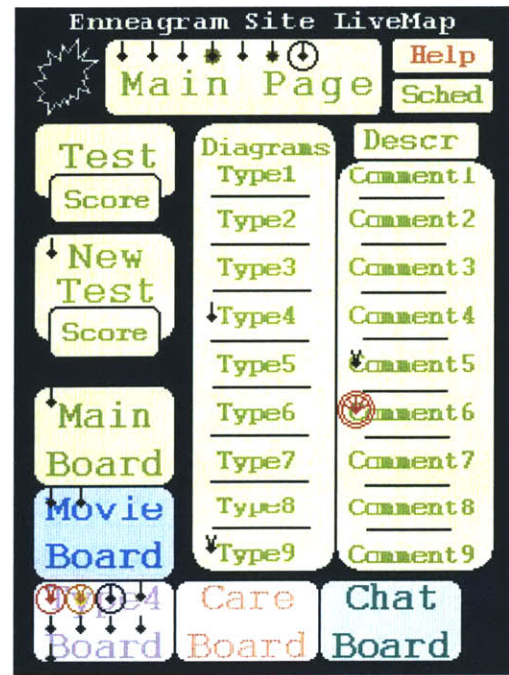d to represent each additional page he visits. At a glance, we can determine how many pages each visitor has seen. We can mouse over a flower to see which pages that visitor has seen. This feature allows users to determine who has seen a page they are interested in and potentially spark a new conversation. Figure 7.4 provides a legend for the different representations of users.

Figure 7.3 The map from 7.2 overlaid with icons for the different users.

visitor	hasn't moved in a while	has been here a while, still active	has been here a while, not active	has been to five different regions of the site	yourself (not a chatter)	chatter	yourself (a chatter)

Figure 7.4 A legend for the different visual representation of users in LiveWeb.

Movement is a useful visual clue to indicate user activities at the site. When a user first visits the site, his icon move from the Internet area, outlined in yellow at the top left corner of the map, to a colored region on the map representing a page on the site. As the user visits more pages, his icon gets more petals and moves from one region of the map to another. The flowers are placed in order, with the most recent flower first, at the top left corner of the region. This means each time a new visitor visits the page, all the existing flowers are moved to the right to make room for the new one. This follows the metaphor of the newcomers jostling the old ones. If a user has not moved

for a certain amount of time, the color of his flower fades. After 5 minutes of no movements, we assume the user is no longer at the site. His icon will then move back to the Internet area, representing the user leaving the Web site.

From looking at the map, in Figure 7.3, one can quickly determine which pages are popular by the number of visitors in the region on the map. The main page is filled with people because that is the page most people visit first at the site. More interesting is the crowd in the region at the lower left corner, which might indicate a new discussion on that message board that has raised a lot of readers' interest.

7.2.3 Interaction

The chat component in LiveWeb allows users on the same Web site to communicate with each other. Users who have selected to chat are asked to enter a name and pick a color for their icon. This gives users some control over how they appear to others. We observe that many chatters indeed see color as a way to express their identity: they tend to stick to a particular color. On the map, chatters are represented with a circle around their flower icons (⊕). All chat messages are broadcast to all the chatters at the site. When a user enters a chat message, his text is displayed above his name on the map. The circle around his icon is also enlarged to make the speaker easier to identify. The circle representation is similar to the one used in ChatCircles [Viegas99].

Our chat facility differs from existing chat rooms in that ours deliberately makes the information at the site the context for conversation. Users may start a conversation based on what they see others are reading at the site. While they talk, they are able to observe changes and happenings at the site. We are interested in allowing interaction among people that's closer to real life: spontaneous interaction born out of people's other activities.

As William Whyte has observed, in real life people tend to block doors and entry to subways as they talk to people, because in these high-traffic areas there is more action and a greater chance of running into someone familiar. Similarly, we think communication among users should not take place in some isolated part of the Web site: users should be able to converse while in full view of the traffic going through the site, and have more options of what to do. The presence display is more than a passive backdrop for interaction. A user can also tell whether other users are following his suggestions of pages to view or see what other users are focusing on.

To create a richer experience, we have also experience with sound. A different sound is played when a user comes to the site or leaves, moves between pages, or sends a chat message. One can easily differentiate the busy versus quiet hours at a site from sound of the movements. This is also useful when users' attention can be focused on browsing, and not on the map itself. Users are able to hear when a new chatter comes or leaves. This scheme has been applied in instant messaging tools like AOL Instant Messenger [AOL99].

7.2.4 Other examples

To illustrate the generality of our approach, we will look at how LiveWeb can be used to show presence on two other sites.

Postcard Site -- This Web site has two main functions: users can select a card to send to a friend or pick a card sent by a friend. We have created a map for the site using colored blocks and circles, as shown in Figure 4. We note that users to the site go through two different paths, which we denoted by two large arrows. The various postcards at the site fall into eight big groups, shown as eight ovals. They make up a large part of the content of the site and the corresponding map.
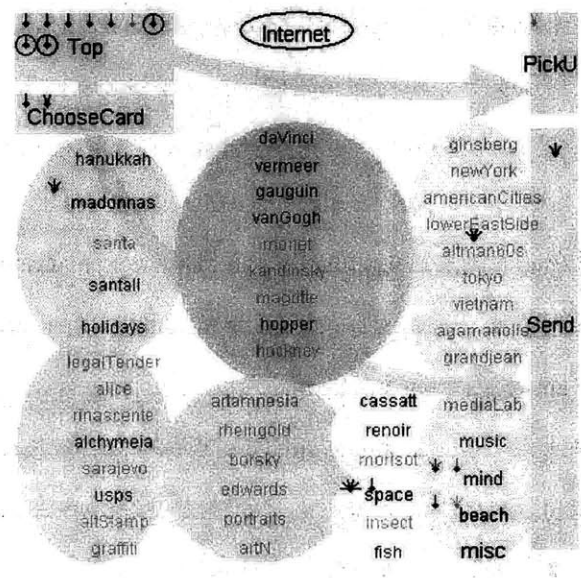
Figure 7.5 An abstract Map of a postcard site showing two paths through the site.

A dynamic presence map can be used to help us understand what people do at the site. In this case, most visitors tend to move from the Top Page (at the top left) down and to the right. So there are quite a few flowers with a single petal starting out at the Top Page. We see that people tend to look through a number of post cards before deciding, noted by a number of flowers with multiple petals in the various postcard circles.

Research Group Site -- This site consists primarily of researchers' homepages, and some research and class informational pages. We feel that the most important aspect of the site are the pages of its various faculty, staff, and

students. Offices and visitors to offices provide an intuitive metaphor for presenting users' movements through the site. We therefore decided to use the floor plan for the group as a site map, as shown in Figure 7.6. In each office, we place the names of its occupants. The remaining pages describing the group are placed in the center lab area, and pages on publications in the remaining unoccupied room to the right. Finally, we add a location for the Internet at the lower right corner of the map to show traffic coming and leaving the site.

It is easy to tell which researcher's page is more popular. For example, Eric's page at the bottom right has a large number of new visitors. He may have some interesting content on his page that is referenced by external sites. So users come in directly to his page. We also show patterns of accumulated user accesses by darkening the text as the homepage it represents gets more visits. This helps users to identify homepages that have been visited more often over time. Currently this feature only applies to map with no hand-painted background, so it did apply to the enneagram site.
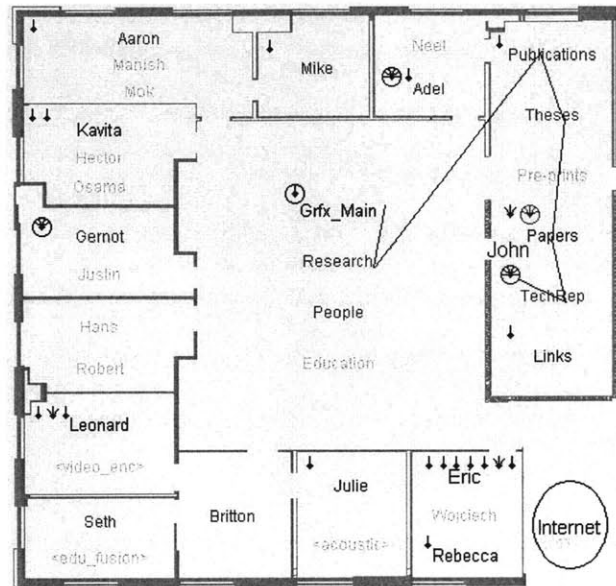
Figure 7.6 A spatially based LiveWeb Map of a research site showing individual's offices.

We have shown three different sample maps for showing presence on Web sites: two based on abstract notions of areas of the site, one based on a physical metaphor. These are hand-crafted to fit the needs of the particular site. In general, a good map should be easy to comprehend following a selected metaphor. A user should be able to tell at a glance which pages are more popular. As he takes time to understand the details of a map, he should also be able to gain more insights about the organization of the site and the movement of other users through the site.

7.3 Related work

A few research projects have dealt with dealt with either user presence display or interaction on the Web. Nelson Minar's crowd visualization shows a loosely clustered set of pages and their animated visitors to give a sense of crowds that have visited a site [Minar99]. This work is closest to ours in conveying a sense of presence at a Web

site using visualization. However, because the pages are laid out in loose clusters, their map is not designed for navigation of the site or for helping users to understand movements within it. It also does not allow interaction among visitors to the Web.

Sociable Web was an early innovation in Web based interaction: a modified Web browser and server system that indicates other individuals on the same Web page and allows the user to strike up conversations or to join in ongoing discussions [Donath94]. It shows promise in combining browsing and surfing. There are also some commercial products for the Web today. Virtual Places, Gooey, and Odigo each use a downloadable browser add-on to allow users to see who is currently visiting the same Web page and to interact with them [Ubique97, Gooey99, Odigo99]. However, because these products are limited to showing presence on the same page, they have not explored the deeper issues of visualizing the crowds and showing the their movements within a site.

In summary, visualizing people at a site and providing means for real-time interaction have not been combined before. LiveWeb aims to create a structured view of the information space, one where users can see others, understand their comings and goings, and interact with them – a closer approximation of the richness of a walk down a real-life street.

7.4 Implementation

LiveWeb is implemented as a client-server system in Java. The user interacts with the system through the client applet from the hosting Web site. The LiveWeb server collects web accesses from both the Web server log and from these Java applets, and serves as a chat server in relaying messages. The applet is divided into two parts: one deals with messages back and forth to the server, and the other the visual interface for users. This modularity allows different visualization schemes to be easily deployed in the future.

This architecture allows us to quickly build a workable prototype that demonstrates our approach. It works with existing browsers. We have easily created the two-frame view that allows users to both browse and interact. By making our clients Java applets, we do not require our users to download or install extra client-side software. All they need is a Java-enabled browser, which over half of the Web users already have in spring 2000, by our count.

However, our chosen implementation does have certain inherent limitations, which should be made explicit. Our system has two types of users: those with the applet, and those without. We can tell which pages they have clicked on the map since the applet can send back the hit. However if a user looks at a different page in the browser, such as by following a link or pressing the back button, the browser may have cached that page already. Since the hit does not go back to the Web server, the LiveWeb server is unable to get the information. This means LiveWeb may not properly display which page a user is currently on. Our log-based approach also means that for a visitor who does not have the applet, we do not know when he leaves the site. Right now, we assume that 5 minutes of no movement means that the visitor has left the site.

Finally, we have implemented LiveWeb as a program that monitors a single log. So when a Web site links to material outside the server, we are unable to display any information for that site. On the map we have created for some sites, we therefore have left out material served from a different site. However, with more time and effort, it is possible to extend LiveWeb to handle multiple servers and combine the log hits from them.

7.5 Future work

So far we have created the LiveWeb system and applied it to three different sample sites. We would like to scale our visualization for larger and more complex sites. This requires not just one map, but perhaps a series of maps for different levels. We would also like to show sites with larger user traffic, so a more scalable user representation is needed. There are two main methods for dealing with the larger traffic: we can filter the users for only those a particular user cares about, or we can use different icons to represent groups of different sizes.

We have used hand-created maps for representing sites. Site maps can also be generated from the site structure. Know what the top-level page point to, and what these pages in turn point to, one can algorithmically generate a map [Chi98]. For example, we can use WebFan to show a hierarchical map of the site. By running a spider through the site, one can easily determine the linkage between pages. This may give a complete but over-complicated structure, lacking emphasis, which is why such maps are used mostly for analysis by site owner, rather than as a navigation aid. One can develop heuristics to pare down the amount of information. One can also use human

intervention to simplify the resulting map. Yet a third way to create site is using data about which links people actually follow, such as in Alan Wexelblat's Footprints [Wexelblat99]. This requires analyzing log accesses, but could create a potentially more useful map for navigation.

Currently we have used a visualization of a site and its traffic as a background for conversation. To further enhance the integration of interaction and browsing, we can add dynamic components synchronized across all the maps. For example, we can specify an area on the map for dynamically showcasing different Web pages from the site. Visitors who have selected that area will see different Web pages being shown in the right hand frame synchronously. This allows multiple users to share an identical view, and comment on the same experience, if they wish. We can also allow users to select pages for displaying to others and get their thoughts on these pages.

We can also encourage users to go the other direction, from chatting to browsing. Based what people are discussing, we can retrieve relevant pages that they can view at their leisure. One simple approach would be a keyword search based on what people have said. Note this is not a way to destroy conversation, but to enrich it. When conversation is intense, these suggestions will likely be ignored. But when the conversation is lagging, people can follow the links for ideas that will further enrich their conversation.

Finally, when more sites have systems like LiveWeb installed, we can allow users to see traffic not only within sites but across sites. Instead of isolated on-line shops, for example, we will be able to see interconnected on-line malls, entertainment centers, and educational complexes, all full of people.

8. Evaluation

In our evaluation, we seek to prove the following two claims:

1. Users prefer being able to see the presence of others on a Web site while chatting.

2. Users prefer being able to see the presence of others on a Web site while surfing.

This is only the first step in verifying our thinking about displaying real-time social presence in on-line information spaces. Many follow-up experiments are needed to verify why people like such displays, what information about others is important to show, and how best to display such information. However, we feel that the most essential piece of our thinking is the assertion that people actually prefer the visual display of presence of other users. Thus, our experiments seek to verify the existence of this preference, by making use of the LiveWeb system described in the last chapter.

8.1 Methods

We evaluated each of the claims with two methods. We conducted experiments on LiveWeb to see whether or not people select the presence display. We also ran on-line surveys to collect user feedback about their selections. Through these two complementary methods, we collected both quantitative data about users' selections, and more qualitative information about what might have led them to make such selections.

8.1.1 Experimental method

As described in the preceding chapter, LiveWeb has been deployed on a public Web site with over 10,000 hits/700 unique users a day. We used this Web site to test how actual users respond to the presence display while chatting and surfing. In our experiments, we presented to users the LiveWeb applet with a switch for turning on or off the presence display (as shown in Figure 8.1); we then record what users select. The naïve measurement would be simply to compare how often a user leaves the presence display on, rather than turning it off. However, most users tend to go with almost any default view, rather than making a selection. To remove the advantage of the default view from our comparison, we presented first the LiveWeb applet showing presence by default, then showing no presence by default. We then compare user selections in these two cases.
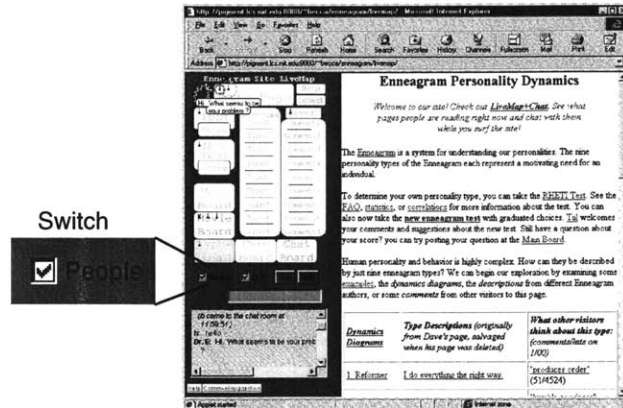
Figure 8.1 LiveWeb Applet with switch for turning presence display on/off.

The data we recorded include users' IP addresses (to identify a user session) and their time-stamped actions, such as turning presence on/off and page views. For the part of our experiment dealing with chat, we recorded in addition what login name a user selects (for more accurate identification across sessions), times of logins/logouts, and the times of chat messages sent.

8.1.2 Survey method

The experimental results could tell us whether people show preference for the presence display, but not why. Our surveys were designed to determine possible reasons for their preferences. We devised our on-line surveys with guidance from Dr. Alberta Lipson. We first conducted a pre-survey, where we asked users several open-ended questions to determine what are appropriate ones and to gain a sense of possible options for answers. This step is important because we did not want to ask our users leading questions or offer false options. The kind of open-ended questions we have asked are as follows:

- Do you like the presence display? If so, why?
- What do you find most useful/interesting?
- What do you find most annoying/frustrating?

Next we created the survey. For each survey question, we offered several choices for the answer, which were formulated on the basis of the pre-survey, and a blank for other possible answers. For each of the pre-surveys and surveys, we added a prominent link from the main page, and encourage user response. Users self-selected to answer these surveys.

8.1.3 Summary of evaluations

We ran three sets of experiments

1. User preference for presence display during chatting

2. User preference for presence display during surfing

3. Same experiment as 2, but on a different site

For each part, we collected data over the course of two weeks. During the first week, the initial view shows presence, and the second week, no presence. In parallel with these experiments, we conducted the corresponding pre-surveys and then the surveys. We now examine the experiment for each part in turn.

8.2 Presence display for chatting

In this experiment, users chose to chat by clicking on a CHAT button from the main page of the site. Then the LiveWeb chat applet would be downloaded. Figure 8.2 shows the LiveWeb chat applet with and without presence display. For one week, users were presented with the presence display (Figure 8.2a) as the default; then in the next week, no presence display as the default (Figure 8.2b). Users could then choose to switch to the other view.
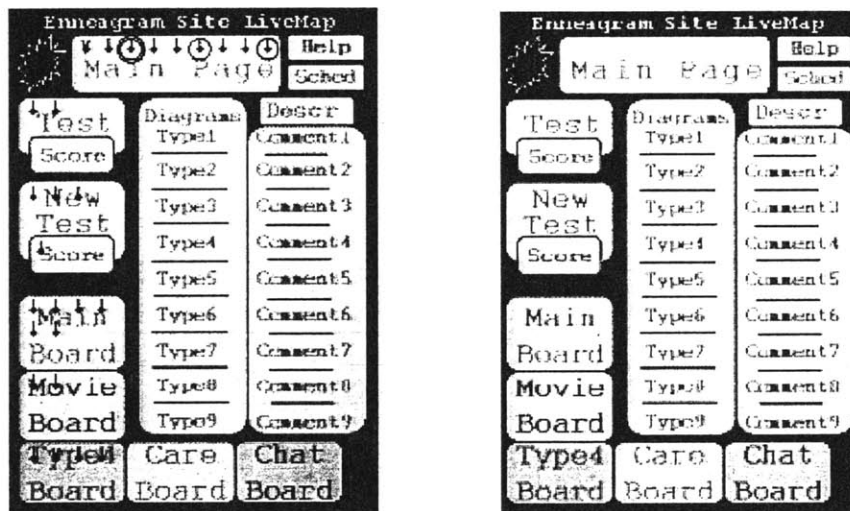
(a) Show presence   (b) No presence

Figure 8.2 The LiveWeb chat applet for chat with presence display turned on (a) or off (b).

8.2.1 Experimental results

We collected valid data for 35 users/115 sessions when presence is displayed by default, and 31 users/179 sessions when no presence is displayed by default. We counted a session from when a user logs in and to the last time he is active, by either sending a chat message or going to a Web page on the site. We did not use the log out time in our calculation, because people often forgot to log out till much later. Sessions that are too short (less than 3 minutes) are considered invalid and ignored.

Table 8.1 summarizes the data collected, showing a strong preference for the presence display. The left column shows results when presence is displayed initially, and the right column when no presence is displayed initially. The rows show a breakdown of people's view selections. If people are indifferent to the presence display, we can expect to see similar percentages in both columns, i.e., that people's decision to switch is independent of the default. If they prefer the presence display, as we proposed, we can expect to see more people switching from default of showing no presence. The latter is indeed what we observe. In the second row we see that the percentage of sessions where the user ultimately switched away from no presence, i.e. switching to presence, is almost eight times greater than those where the user ultimately switched away from presence. In fact, since we have colored the percentages according to what view the users will see as a result of their selection, we can see that the numbers in red, for showing presence, are consistently larger than the number in blue, for no presence view, in each row.

	Initial: Show presence	*Initial: No presence*
Out of total sessions:		
Kept initial view (never switched):	88%	44%
Switched and kept new view:	7%	54%
Switched and returned to initial view:	5%	2%
Out of sessions where the user switched:		
Kept new view:	57%	96%
Returned to initial view:	43%	4%

Table 8.1 A summary of results from the experiment on user preference for presence display during chatting. Numbers are colored according to which view the user sees as a result of their selection: red indicates showing presence, blue no presence.

We can also look more closely at the last two rows of Table 8.1 to see what happens when people attempted to switch. The left column shows that in sessions where a user initially turned the presence display off, almost half result in the user later changed his mind and turned it back on. In the right column, we see that users who initially turned on the presence display overwhelmingly decided to keep it on. The results show people who turned on the presence display were much more likely to stick with their decision, showing again a greater preference for the display.

Since the users are required to enter a login each time, we can also trace a user's decision across sessions by their logins. This helps us to gain a clearer idea of individual users' preference. Table 8.2 shows a parallel of the results presented in Table 8.1, but this time summed by users. At a glance, we see that the red numbers (for resulting view that shows presence) are again greater in each row. The second row of Table 8.2 shows that in fact a large percentage of users tried both views. The percentage of people who tried to switch on presence display (right column) remains much larger than those who tried to switch it off (left column).

	Initial: Show presence	Initial: No presence
Out of total users:		
Always kept initial view (never switched):	71%	16%
Switched at least once:	29%	84%
Out of users who switched:		
Kept new view more often:	50%	100%
Returned to initial view more often:	50%	0%

Table 8.2 A summary of results from the experiment on user preference for presence display during chatting, this time counted by user. Numbers are colored according to which view the user sees as a result of their selection: red indicates showing presence, blue no presence.

The last two rows of Table 8.2 show out of those who switched, how many decided to switch to keep the new view more often. We see that people were equally split when the default showed presence (left column). But when the default did not show presence, not only did far more people try switching but that everyone who had switched chose to keep the presence display on more often (right column). This is indeed encouraging.

8.2.2 Survey results

So why did people prefer the presence display? Our survey revealed some possible reasons. We collected 13 survey results back from the users, out of which 12 reported liking the presence display. The top reasons for preferring the presence display were that it

- Is interesting to watch
- Shows where a lot of people are
- Indicates new post to the message boards
- Shows changes to traffic

The results show that users found the presence display interesting in itself, but also saw it as more than just a novelty. People used it to gain information about what others at the site are doing. The main reason for disliking preference display was that it can be distracting.

8.3 Presence display for surfing

We repeated the same experiment for users' preference for presence display when surfing the Web site. Users were again initially presented with one of the two views shown in Figure 8.3, and given the option to switch to the other. In week one, users were initially shown presence (Figure 8.3a), and week two, no presence (Figure 8.3b).

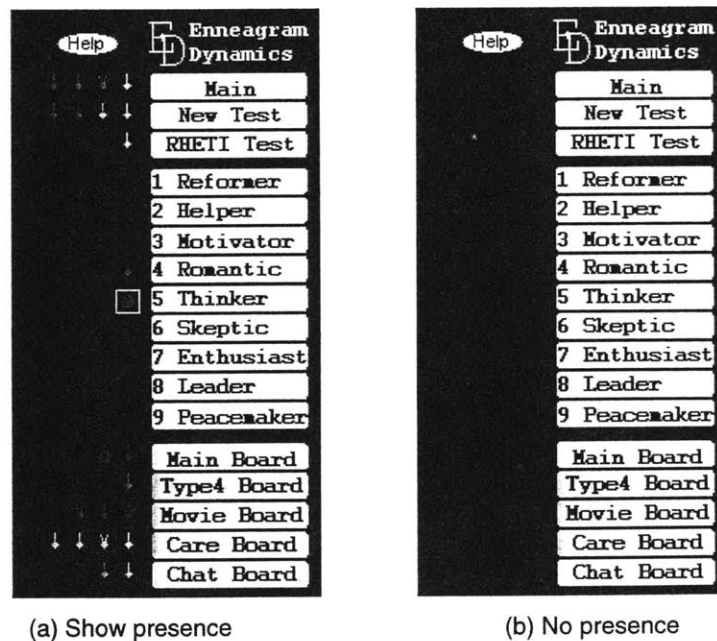(a) Show presence   (b) No presence

Figure 8.3 LiveWeb Applet for surfing with presence display turned on (a) and off (b).

In the chatting experiment, users needed to explicitly choose to chat. In this experiment, the LiveWeb applet would automatically be loaded when users first came to the Web site, so all users with the appropriate browser saw the applet. We did not ask the surfers to log in, because we wanted to get a wider population of users for the applet and felt that asking them to log in may discourage some from it. Also, there was less need to identify individuals since they could not chat. However, this meant we could only track individual sessions, but not users across sessions.

8.3.1 Experimental Results

We examined all user sessions lasting longer than three minutes: 463 sessions for default on, and 552 sessions for default off. The results are shown in Table 8.3. We see that the percentage of people who switched from default is much less than in the previous experiment. The difference was that in this experiment, we automatically showed all surfers with the available browser the LiveWeb applet, whereas in the last one users have to click a button to see the applet. This meant we got a much larger and diverse audience, but of people who only expected to see a Web page.

	Initial: Show presence	Initial: No presence
Out of total sessions:		
Kept initial view (never switched):	92%	82%
Switched and kept new view:	5%	15%
Switched and returned to initial view:	3%	2%
Out of sessions where the user switched:		
Kept new view:	59%	88%
Returned to initial view :	41%	12%

Table 8.3 Results from the experiment on user preference for presence display during surfing. Numbers are colored according to which view the user sees as a result of their selection: red indicates showing presence, blue no presence.

That aside, we see the same trend as in the chatting case, that users show a greater preference for presence display. Row two shows that the users turned on the presence display three times as often as the opposite. The last two rows of Table 8.3 show of attempts to switch, how many resulted in the user switching back to the default. We see that users returned to the presence display (left column) over three times as often as to the no presence display (right

column). The difference here is smaller than in the last case. We speculate that the chatters might be more tuned in to the social side of the Web than the surfers, and therefore adopted the presence display in greater percentages.

8.3.2 Survey results

For our survey, we collected 30 results from users: 22, or 73%, like the presence display. The top reasons for preferring presence display are that it

- Is fun to see what others are doing
- Is interesting to watch
- Provides a sense of connection to others
- Shows how many people are on
- Shows where are a lot of people

Most of the reasons are similar to those given in the last section for the chatting case. Eight people, or 27%, did not like the presence display. Half of them claimed that it is too distracting. The other half were split: some were indifferent, some found it invaded privacy. One person in particular voiced a very strong opinion against the presence display: she felt she didn't care to know what others are doing and neither should other people know about her activities.

8.4 Presence display for surfing – second site

We conducted the previous two rounds of evaluation on the same Web site. To ensure that the findings are not unique, we also reran the surfing experiment on a different site, the MIT Computer Graphics Group Site. The set-up in this case was slightly different: users must choose to view the LiveWeb applet by clicking on a link from the main page. Then the applet would be loaded, and the users once again presented initially with one of the two views shown in Figure 8.4, and allowed to switch to the other.

Note that we have also used a different representation for people here (🕴🕴): the humanoid icons get bigger as user visits more places. One reason for the change was that we want to see whether users will still show preference for presence display with a slightly different interface. A second was that some people have reported that the flower metaphor for representing users is at first hard to grasp at first, so we wanted to experiment with something that more closely suggests a user.
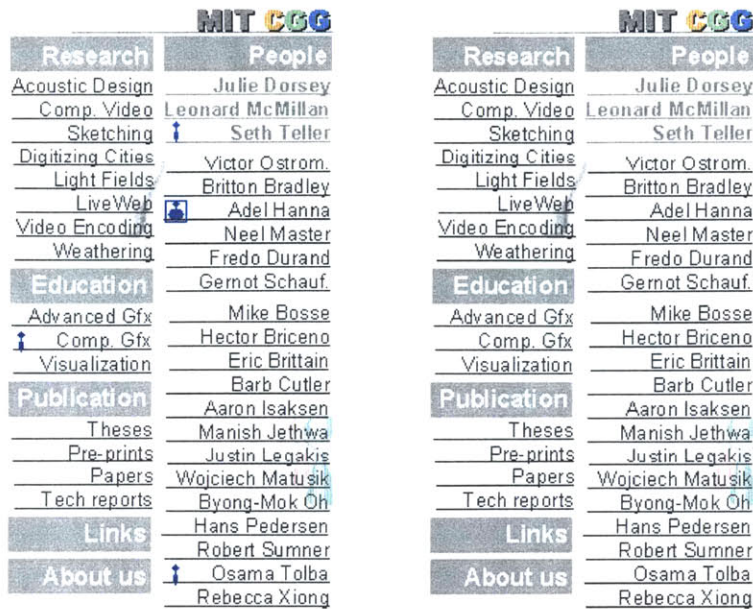
(a) Show presence   (b) No presence

Figure 8.4 LiveWeb Applet for surfing on a second site with presence display turned on (a) and off (b).

8.4.1 Experimental Results

We examine all the user sessions: 96 sessions for default on, and 92 sessions for default off. The results are shown in Table 8.4. We see the same trend as in the surfing experiment conducted on the first site, that users show a greater preference for presence display than otherwise. The different icons used to represent people also had not changed the overall findings.

	Initial: Show presence	Initial: No presence
Out of total sessions:		
Kept initial view (never switched):	86%	64%
Switched and kept new view:	0%	35%
Switched and returned to initial view:	14%	1%
Out of sessions where the user switched:		
Kept new view:	0%	97%
Returned to initial view:	100%	3%

Table 8.4 Results from the second experiment on user preference for presence display during surfing. Numbers are colored according to which view the user sees as a result of their selection: red indicates showing presence, blue no presence.

However, the preference for presence shown here is much more prominent. The second row of Table 8.4 shows that nobody switched from the initial presence display (left column), while roughly a third overall switched from the initial no presence display (right column). The last two rows of Table 8.4 show, of the times when user switched, how many resulted in the user switching back to the default. The numbers here are again overwhelming: everyone initially presented with the presence display returned to it (left column); and almost no one initially presented with no presence display returned (right column). We speculate that this strong preference is due to the fact that the visitors on the computer graphics page were more technically savvy than those on the enneagram site, and therefore more easily understood the presence display and found it useful. Also, users self-selected to see the applet, unlike in the surfing experiment in the last experiment.

Because of the small sample size, we did not collect enough pre-survey result to run a survey. However, six out of the seven pre-survey results did show preference for the presence display. Reasons sited for the preference were to see others' browsing patterns, to find out where others congregate, to be drawn to things that one would otherwise not notice. The first two reasons are by now familiar; the third reason suggests another application of social navigation. Three people also independently suggested that having a chat function would be most helpful.

8.4 Discussion

Our experiments show that during both chatting and surfing, users preferred to see the display of presence. They clicked to turn on the presence display more often, and they returned to the presence display more often than otherwise. In the chatting experiment, users showed an overwhelming preference for the presence display. For the surfing experiment conducted on the same site, the preference was still significant, but to a lesser degree. One reason may be that whereas chatters must choose to see the applet, all surfers with the appropriate browser saw the applet automatically. The chatters were also more likely to be interested in the social rather than the purely informational side of the Web experience.

Running the surfing experiment on a second site gave us similar results. Users again chose to turn on the presence display more and return to the presence display more. The number of people who press the button is greater than in the surfing experiment on the first site, possibly because people self-selected to see the presence display, and

because of visitors to the second site (on a computer graphics research group) may have greater technical knowledge.

In surveys, people cited various reasons for preferring the presence display. They found it interesting to watch the dynamic movement of icons showing activities of other users at the site, although a small number found it distracting. We believe that there is a real audience for the kind of presence display we have developed. Many users quickly catch on to the display, for its ability to convey a sense of social connection and to show what others find interesting about the site. Others may be late adopters who accept new technology more gradually. There are also those for whom the Web experience is about static information only, and therefore will not find this display useful.

User comments provide useful feedback for the future re-design of our system. The comments from the surfers on the second site suggest that users like the ability to chat with other uses, a function we should not only keep but expand. We also found that privacy is a concern to a small but vocal minority, and should be addressed appropriately. While the vast majority of people do not mind being tracked when not personally identified, a small number dislike any sense of being tracked. We should therefore allow users to opt out of the system, and to not have their activity shown. In return, these users should not be allowed to see others' activities, following the exchange of information scheme described in Section 2.6.

We might want to allow users to create a profile that specifies whether they want to be tracked. The profile can also specify whether they want presence display or not, since a large number of people tend to go with the default view, and do not bother to switch even though they may have a greater preference for the other view.

9. Conclusion

Human beings are social animals. In real life, we are always in contact with other people, whose activities help us decide where to go, what to see, and whom to talk to. We have proposed that users would also prefer to be able to see the presence of others in virtual environments like the Web. This awareness of others could potentially help them to navigate socially, gain a greater sense of connection to others, and make use of the information content available to enrich their interactions. We have created the LiveWeb system to display real-time activities of others on Web sites; we have installed it on two Web sites and conducted evaluations of the system in action.

Our experiments using the LiveWeb system indicated a strong user preference for the display of presence. During both chatting and surfing, and across two separate sites, people consistently selected the presence view and returned to it much more often than the no presence view. Most people have correctly interpreted our visualization as not merely a set of moving icons, but representations of real people and their activities: the top reason for preferring the presence display was to see what others are doing at the site. Moreover, people reported that they actively use the display to gain information about their fellow visitors: they observe how many users are at the site, and how this traffic changes over time.

Although we have not conducted formal studies of the cause of users' preference for presence display, our survey results provide informal evidence supporting the three functions of presence we proposed earlier:

- Social navigation -- Many people have used the presence display to see which page on the site attracted a large number of others, which is an important step in social navigation. Some respondents further indicated that they checked the display for new crowds that might indicate a new post to the board.

- Social connection -- Roughly one third of the users who prefer presence display reported that the display helped them gain a greater sense of connection to others. Moreover, awareness of others can remind people of the opportunity for social interaction. In our surfing-only version of the LiveWeb interface, for example, several users independently asked for the ability to chat with others at site, whose presence they are now aware of.

- Interaction with information as a context -- Half of the users who preferred the presence display think that it has affected how they chat with others. Some used it to check what others are looking at or even to comment on it. These reactions suggest that our visualization can potentially be used to gain a better understanding of the interest of others and to create more effective openings for conversation. More analysis of the chat text is needed to see what people have actually talked about, and the degree to which the informational content of the site has played a part.

We feel our results are encouraging and suggest future studies to quantify how the awareness of others can affect people's on-line behaviors.

If we are persuaded that the people find the presence display to be useful, we can then ask how to design effective visualizations of presence. Our experience with LiveWeb and other designs suggest four design considerations:

- Visual simplicity

- Visual variety

- Motion

- User exploration

Visual simplicity requires no-frills designs and intuitive visual encodings. The blocky shapes of the nodes in NetViz consumed too much screen space and user focus, whereas the lean lines of WebFan offered a much cleaner overall look. The yellow-green-blue color encoding in NetViz generated too much contrast and was difficult to remember, whereas the use of fading to represent older messages in PeopleFlower offered a more intuitive encoding.

Visual variety entails combining several encodings to convey more information about the objects represented. With this technique, we can create a virtual crowd that is varied and interesting, like those in real life. In LiveWeb, we started with the PeopleFlower metaphor and used the number of petals and their brightness to encode other cues about user behaviors. While most people were eventually able to comprehend the differences between the different icons in LiveWeb, these icons initially mystified some users. We therefore experimented with simpler humanoid icons. Though these graphical representations of people were easier to understand, they were less capable of

conveying different cues legibly. While it is desirable to maximize both visual simplicity and variety, we found in practice that we may often need to trade off between the two. The placement of emphasis depends on whether the goal of the visualization is easy comprehensibility or information-richness.

We found that motion can convey information in a compelling manner: people were often instinctively drawn to things that move. The falling metaphor used in NetViz could indicate clearly when certain nodes were filtered out. Moving icons in LiveWeb were essential for conveying a sense of others' activities within the site. We speculate that the its dynamic nature was one of reasons that people found the LiveWeb display interesting to watch. However, motion could also have the negative effect of creating distraction, as was pointed out by a few users. We may want to allow users to choose between a more engaging display of motion and another, less dynamic, display of presence.

Finally, we might want to decide what important information should displayed initially, and what details should be withheld. This would allow a cleaner visualization at first and would then invite users to interact and learn more about what they find interesting. We advocate responsive displays that enable user exploration. WebFan, for example, presents users with an instant overview of the structure of a conversation, and then expands the threads of interest to show more details as a user interacts with the visualization. LiveWeb allows users to focus on representations of specific individuals and determine what paths they took through the site.

We have discussed the theory, designs, and design principles that have culminated in the LiveWeb, our system for visualizing live user presence. Actual use of the system has showed that users strongly prefer the presence display, and has offered hopeful glimpses of its potential benefits. We believe that allowing users to become aware of the activity of others can help to move the Web beyond a global information repository and toward a global interaction space.

References

Real people, real environments

[Alexander65] Alexander, C., "A City Is Not a Tree", *Architecture Forum* 1965, in *The City Reader*, R. LeGates and F. Stout, eds. Routledge. 1996.

[Aronson54] Aronson, E., *The Social Animal*, W. H. Freeman and Company, 1954.

[Blumer86] Blumer, H., *Symbolic Interactionism: Perspective and Method*, University of California Press, March 1986.

[Brand94] Brand, S., *How Buildings Learn: What Happens After They're Built*, New York, NY, Penguin, 1994.

[Goffman59] Goffman, E., *The Presentation of Self in Everyday Life*. New York: Doubleday, 1959.

[Jacobs61] Jacobs, J., *The Death and Life of Great American Cities*, 1961.

[James71] James, M. et al, *Born to Win : Transactional Analysis With Gestalt Experiments*, Perseus, 1971.

[Hardin68] Hardin, G. "The Tragedy of the Commons", *Science* 162: 1243-48, 1968.

[Lynch60] Lynch, K., *The Image of the City*. Cambridge, MA: MIT Press, 1960.

[Lynch72] Lynch, K., *What Time Is This Place?*, Cambridge, MA: MIT Press, 1972.

[Milgram67] Milgram, S. "The Small World Problem", *Psychology Today*, 1 1967.

[Milgram77] Milgram, S.. *The Individual in a Social World*. Reading, MA: Addison-Wesley, 1977.

[Moreno34] Moreno, J. *Who Shall Survive*, New York: Beacon Press, 1934.

[Oldenburg89] Ray Oldenburg, *The Great Good Place: Cafes, Coffee Shops, Community Centers, Beauty Parlors, General Stores, Bars, Hangouts, and How They Get You Through the Day* (New York: Paragon House, 1989).

[Raskin60] Raskin, E., "On theNature of Variety", *Columbia University Forum*, 1960, quoted in [Jacobs61].

[Wellman79] Wellman, B., "The Community Question: The Intimate Networks of East Yorkers", *American Journal of Sociology*, 84, 1979.

[Whyte88] Whyte, W. , *City: Rediscovering the Center*, 1988.

[Wirth38] Wirth, L. , *Urbanism as a way of life*, 1938.

Virtual environments

[AlphaWorld99] AlphaWorld Web site, http://www.alphaworld.com, 1999.

[ACM99] ACM Digital Library Web site, http://www.acm.org/dl/, 1999.

[Biocca95] Biocca, F. and M.R. Levy, *Communication in the age of virtual reality*, Hillsdale, NJ: Lawrence Erlbaum Associates, 1995.6

[Biocca97] Biocca, F., "The Cyborg's Dilemma: Progressive Embodiment in Virtual Environments", *Journal of Computer Mediated Communication* 3 (2), September 1997.

[Bly93] Bly, S., S. R.Harrison, and S. Irwin, Media spaces bringing people together in a video, audio,and computing environment; *Commun. ACM* 36, 1 (Jan. 1993), Pages 28 -46.

[Borovoy98] Borovoy, R., Martin, F., Vemuri, S. Resnick, M., Silverman, B. and Hancock, C., "Meme tags and community mirrors: Moving from conferences to collaboration", *ACM Conference on Computer Supported Collaborative Work*, 1998.

[Chandler99]   Chandler, D., Personal Home Pages and the Construction of Identities on the Web, URL http://www.aber.ac.uk/~dgc/webident.html, 1999.

[Heeter92]   Heeter, C., "Being There: The subjective experience of presence", *Presence: Teleoperators and Virtual Environments*, MIT Press, 1992.

[Kollock96]   Kollock, P and Smith, M. "Managing the Virtual Commons: Cooperation and Conflict in Computer Communities", In *Computer Mediated Communication: Linguistic, Social and Cross-Cultural Perspectives*, Susan Herring, ed. John Benjamins, 1996.

[Lombard97]   Lombard, M. and T. Ditton, "At the Heart of It All: The Concept of Presence", *Journal of Computer Mediated Communication* 3 (2), September 1997.

[Rheingold93]   Rheingold, Howard. 1993. *The Virtual Community: Homesteading on the Electronic Frontier.* Reading, MA: Addison-Wesley Pub. Co.

[Reid 94]   Reid, E., *Cultural formations in text-based virtual realities.* Thesis, Dept. of English, University of Melbourne, 1994.

[Rice93]   Rice, R. E., Media appropriateness: Using social presence theory to compare traditional and new organizational media. *Human Communication Research, 19(4),* 451-484, 1993.

[Short76]   Short, J., E. Williams, B. Christie, *The social psychology of telecommunications,* London: Wiley, 1976.

[Smith97]   Smith, M., Netscan: Measuring and Mapping the Social Structure of Usenet, *17th Annual International Sunbelt Social Network Conference,* San Diego, California, February 13-16, 1997.

[Smith99]   Smith, M., "Invisible Crowds in Cyberspace: Measuring and Mapping the Usenet", *Communities in Cyberspace* , Routledge, 1999.

[Sproull95]   Sproull, L. and S. Kiesler, *Connections: New ways of working in the networked organization,* MIT Press, 1995.

[Turkle95]   Turkle, S., *Life on the Screen: Identity in the Age of the Internet.* NY: Simon & Schuster, 1995.

[Whittaker98]   Whittaker, S., Terveen, L., Hill, W., Cherny, L., "The Dynamics of Mass Interaction", *ACM Conference on Computer Supported Collaborative Work,* 1998.

Visual interface design

[Ackerman95]   Ackerman, M. and B. Starr, Social activity indicators: interface components for CSCW systems; Proceedings of the 8th ACM symposium on User interface and software technology , 1995.

[Ahlberg94]   Ahlberg, C. and Shneiderman, B.,Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Human Factors in Computing Systems CHI '94 Conference Proceedings,* 1994.

[Arnheim69]   Arnheim, R., *Visual Thinking,* University of California Press, 1969.

[Art97]   *The Art Book,* London: Phaidon Press, 1997.

[Bailey93]   Bailey, G., "Iterative methodology and designer training in human-computer interface design", CHI '93.

[Benford95]   Benford, S., J. Bowers, L.E. Fahlén, C. Greenhalgh, and D., Snowdon, User embodiment in collaborative virtual environments, Proceedings of CHI '95.

[Benford97]   Benford, S., C. Greenhalgh and D. Lloyd, "Crowded collaborative virtual environments", CHI '97 Conference on Human Factors in Computing Systems, 1997.

[Brilliant93]   Brilliant, R., *Portraiture,* Harvard University Press, 1993.

[Butler96]   Butler, K., "Usability engineering turns 10", Interactions, 3, 1 (Jan. 1996), Pages 59-75.

[Card99]   Card, S.K., J.D. MacKinlay (ed.), B., Shnieiderman (ed.), *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufman, 1999.

[Chi98]   Chi, E., et al, "Visualizing the Evolution of Web Ecologies", CHI '98 Conference on Human Factors in Computing Systems, 1998.

[Damer97]   Damer, B., Avatars! : Exploring and Building Virtual Worlds on the Internet, Peachpit Press, 1997.

[Donath95]   Donath, J., "Visual Who: Animating the affinities and activities of an electronic community", at ACM Multimedia 95, Nov 5-9, San Francisco, CA.

[Donath95b]   Donath, J., The Illustrated Conversation,Multimedia Tools and Applications. Vol. 1, March 1995.

[Donath96]   Donath, J., Inhabiting the Virtual City. Ph.D. dissertation, MIT Media Lab, 1996.

[Donath99]   Donath, J., "Identity and Deception in the Virtual Community", Communities in Cyberspace, Kollock, P. and Smith M. (eds)., London: Routledge,99.

[Donath99b]   Donath, J., K. Karahalios, and F. Viegas, "Visualizing conversations". In Proceedings of HICSS 32, January 5-8, 1999.

[Dourish92]   Dourish, P. and V. Belotti, Awareness and Coordination in Shared Workspace, CSCW '92.

[Foltz98]   Foltz, M., "Designing navigable information spaces", Masters Thesis, MIT Lab for Artificial Intelligence, 1998.

[Harris97]   Harris, R. L., *Information Graphics: A Comprehensive Illustrated Reference: Visual Tools for Analyzing, Managing, and Communicating*, Management Graphics, 1997.

[Hill92]   Hill, W.C., Hollan, J. D., Wroblewski, D. and McCandless, T. 1992. Edit wear and read wear. Proceedings of CHI '92. 3-9

[Hill93]   Hill, W.C. and J. D. Hollan, History-Enriched Digital Objects, CFP'93.

[InXight99]   InXight Web site http://www.inxight.com,1999.

[Jeffries91] Jeffries, R., J.R. Miller, C. Warton and K. Uyeda, "User interface evaluation in the real world: a comparison of four techniques" , CHI '91.

[Kahn99]   Kahn, P., "Mapping Web Sites: Planning Diagrams to Site Maps", Dynamic Diagrams, Inc., http://www.dynamicdiagrams.com/seminars/mapping/maptoc.htm, 1999.

[Kautz98]   Kautz, H and Selman, B., "Creating Models of Real-World Communities with ReferralWeb", Working notes of the Workshop on Recommender Systems, held in conjunction with AAAI-98, Madison, WI, 1998.

[Krebs96]   Krebs, V., Visualizing Human Networks, Release 1.0, 12 February 1996.

[Kurlander96] Kurlander, D., T. Skelly, and D. Salesin, Comic Chat, Proceedings of SIGGRAPH '96.

[Jeffery97]   Jeffery, C., "A menagerie of program visualization techniques", in Software Visualization: Programming as a Multimedia Experience, MIT Press, 1997. [Headcount98]   Headcount.com, http://www.headcount.com,1998

[Lamping95]   Lamping J., Rao R., Pirolli R., "A Focus and Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies", CHI '95 Conference on Human Factors in Computing Systems, 1995.

[Mukherjea95] Mukherjea, S. and Foley, J. D., "Visualizing the World Wide Web with the Navigational View Builder", WWW '95.

[Nielsen90]   Nielsen, J. and R. Molich, "Heuristic evaluation of user interfaces", CHI '90.

[Norman90]   Norman, D., The Design of Everyday Things, New York: Doubleday, 1990.

[Pedersen97] Pedersen, E.R. and T. Sokoler, "AROMA: abstract representation of presence supporting mutual awareness", CHI '97.

[Plaisant96] Plaisant, C. et al, "LifeLines: Visualizing Personal Histories", CHI '96 Conference on Human Factors in Computing Systems, 1996.

[Rankin98]   Rankin, P.J., et al, "StarCursors in ContentSpace: Abstraction of People and Places", Workshop on Personalized and Social Navigation in Information Space, 1998.

[Schutte98]   Schutte, A., Patina: layering a history-of-use on digital objects, S.M. Thesis, MIT Media Lab, 1998.

[Terveen98]   Terveen, L., and W. Hill, Finding and Visualizing Inter-site Clan Graphs, Proceedings of CHI '98.

[Tufte83]   Tufte, E. R., 1983. The Visual Display of Quantitative Information. Cheshire, CT: Graphics Press.

[Tufte90]   Tufte, E.R.,-. 1990. Envisioning Information. Cheshire, CT: Graphics Press.

[Viegas99]   Viegas, F., and J. Donath, "Chat Circles", in Proceedings of CHI '99.

[Weiser91]   Weiser, M.,"The Computer of the Twenty-First Century", in Scientific American, 10, September 1991.

[Weiser96]   Weiser, M., and J.S. Brown Designing Calm Technology, PowerGrid Journal, v 1.01, July 1996.

[Wexelblat99]   Wexelblat, A., "History-based tools for navigation", Hawaii International Conference on System Sciences, Jan 5-8, 1999, IEEE Press.

[Whittaker94]   Whittaker, S., Frohlich D., Daly-Jones, O., "Workplace Communication,: What is it like and how might we support it?", CHI '94 Conference on Human Factors in Computing Systems, 1994.

[Whittaker97]   Whittaker S, Swanson J, Kucan J, Sidner C. "TeleNotes: managing lightweight interactions in the desktop", ACM Transactions on Computer-Human Interaction, vol.4, no.2, June 1997, pp.137-68. Publisher: ACM, USA.

Other

[AOL99]   American Online, Inc. Web Site, http://www.aol.com/, 1999.

[Amazon99]   Amazon.com Inc. Web Site, http://www.amazon.com, 1999.

[Deja99]   Deja News Web Site, http://www.deja.com/, 1999.

[EachMovie98]   EachMovie Web Site, http://www.eachmovie.com/, 1998.

[Gooey99]   Gooey Web Site http://www.gooey.com/, 1999.

[Goldberg92]   Goldberg, D., B. Oki, B., D. Nichols, D.B. Terry, Using Collaborative Filtering to Weave an Information Tapestry. Communications of the ACM, December 1992, Vol 35, No12, pp. 61-70.

[Kowalski99]   Kowalski, M., Markosian, L., et al, Art-Based Rendering of Fur, Grass, and Trees, in SIGGRAPH '99.

[Raskar99]   Raskar, R. and M. Cohen, "Image Precision Silhouette Edges", Symposium on Interactive 3D Graphics, Atlanta GA, April 1999.

[Resnick94]   Resnick, P., I. Neophytos, I., S. Mitesh, P. Bergstrom, and J. Riedl, GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In Proc. of CSCW '94: Conference on Computer Supported Cooperative Work (Chapel Hill, 1994), Addison-Wesley.

[ThirdVoice99] ThirdVoice Web site http://www.thirdvoice.com/, 1999.

[Ubique97]   Ubique Web site http://www.ubique.com/, 1997.