Browsers to Support Awareness and Social Interaction

Alison Lee
IBM T.J. Watson Research Center

Andreas Girgensohn
FX Palo Alto Laboratory

Jun Zhang
University of Michigan

Browsers that combine social visualizations and tools let newcomers and visitors explore information and patterns. Here we present social browsers for two Web communities.

Information sharing and social interaction are the Web's main features that have enabled online communities to abound and flourish. However, the Web is lacking cues and browsing mechanisms for the online social spaces. That is, such spaces look and feel abstract and informational rather than inviting for social interaction. Consequently, newcomers and occasional visitors to such sites cannot quickly assess who the participants are; what the purpose, tenor, and norms of the space are; which sorts of newcomers are welcomed; and what forms of participation are valued. Thus, users need support for social browsing in such spaces.

On the flip side, active contributors to social spaces (such as Web communities) lack the means to present a public face to visitors. Gathering places in the physical world, in contrast, reflect artifacts, patterns, and traces of social presence, activities, and organization.[1] They convey the life, character, relationships, and social cues that are important for social interaction and organization. Such social information and patterns exist in the electronic world but are buried in the content-centric information. (For more information see the "Social Visualizations" sidebar.)

The challenge of creating social browsing tools to access such social information and patterns is of interest as a visual analytic problem for two reasons. First, unlike traditional users of typical visual analytic tools, Web community visitors and participants are neither necessarily literate on visualizations, nor do they view their social inquiry activities as explicit forms of visual analysis activities. Hence, these users of social browsing require tools that focus not only on information visualization but also on visual design and user support mechanisms.

Second, to apply visualization techniques to the domain of social computing and the problem of providing cues and conveying social dynamics in online communities, we would apply visual analytics to the

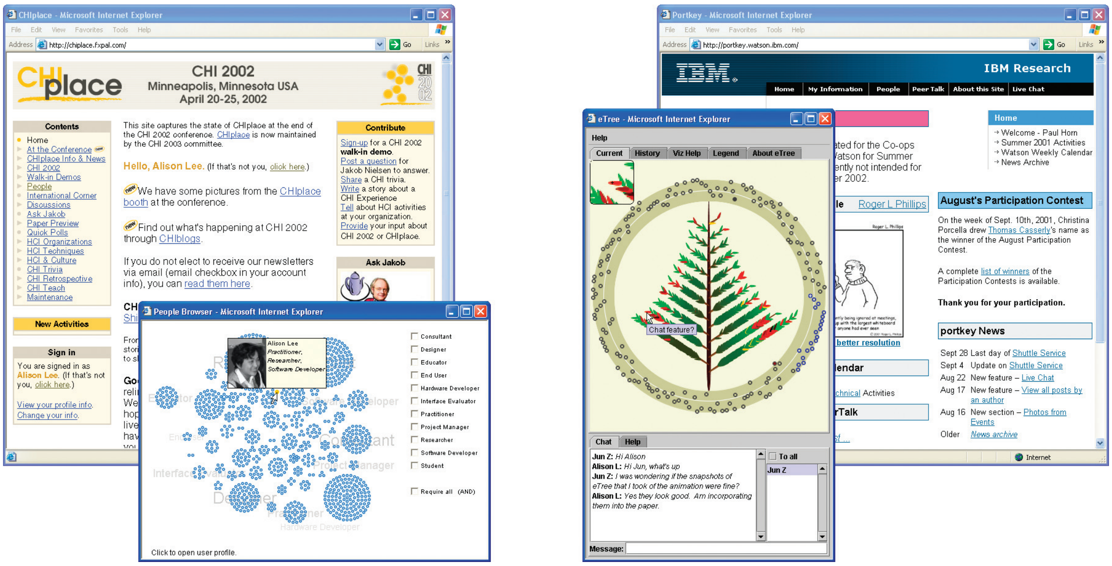

1 (a) CHIplace pictured with the People browser and (b) Portkey pictured with the eTree browser.

Published by the IEEE Computer Society

properties and data representing human interaction. Thus, the social visualization and browsing problem presents an interesting variation of visual analytics.

We demonstrate, using two case study examples, how social browsers exploit visual design and information visualization techniques to present social information to nonexpert visualization users. Also, we argue that such browsers should incorporate tools that let newcomers and visitors view and explore the evolution of information and the emergence of patterns in online social spaces.

The browsers

The case studies use two social browsers (People and eTree) that we developed for two different community Web sites—CHIplace and Portkey (see Figure 1). Each browser visualizes a particular community property or characteristic of human interaction over time and supports exploration and discovery of the social context information for its Web site.

These two social browsers feature several innovative aspects. First, each visualization presents a portrait of a facet of a community derived from different social information. The CHIplace People browser provides a community map of the site members' different roles. The Portkey eTree browser portrays the participants and the extent of discussion participation through a visual ecosystem metaphor.

Second, both visualizations evolve and dynamically adjust to reflect the activities in the community space. As such, the visualizations make explicit certain social information that lets visitors and participants observe—over time—patterns, human presence, and/or social activities and organizations occurring on the Web sites.

Third, the browsers incorporate several visual design and information visualization techniques to create an intuitive and engaging environment for exploring the social activity of the Web communities.

Finally, the browsers' visualizations have supplemental functionality (such as dynamic query) and information visualization techniques (such as animation, overview, and detail). These additional capabilities aid users in their exploration and discovery of the social information occurring in the community Web sites.

Social Visualizations

Several related efforts in awareness systems and social visualization have influenced our work. Extensive work has been done in the awareness area in computer-supported cooperative work to support the formal and informal interactions within distributed groups.[1] Researchers have created tools that provide information about the context of a work environment, letting the members maintain awareness and guide their own actions. While many early efforts explored the use of audio and video to keep abreast of a work environment, subsequent efforts have focused on identifying important social cues and providing lightweight, low-cost, minimalist techniques to convey such cues. These efforts used icons, graphics, and other forms of indicators such as color light-emitting diodes to indicate people's availability state or graphical marks on edges of documents to show how extensively they were viewed.[2,3] One aim of our work is to support social interaction and awareness. We achieve this by providing important social cues such as presence of people, availability of people, or information about members of the online community.

Unlike some of the work on the visualization of social data such as Lifelines,[4] much of the current social visualization work focuses on conveying information about the online world and its participants. Systems such as Conversation Map, Netscan, and Personal Map let users explore social information in small- and large-scale conversation spaces.[5-7] These social visualizations were used typically as analysis and study tools for finding important patterns in data. More recently, Donath[8] has suggested that social visualizations are valuable in giving the participants a better grasp of their online social space. Examples include Comic Chat[9] and Chat Circle.[10] This latter research is similar to the work in awareness systems and the work presented here. Specifically, we are concerned with developing visual design techniques for creating simple, legible, and evocative visualizations.

References

1. O. Liechti and Y. Sumi, "Editorial: Awareness and the WWW," *Intl J. Human-Computer Studies*, vol. 56, no. 1, 2002, pp. 1-5.
2. M.S. Ackerman and B. Starr, "Social Activity Indicators: Interface Components for CSCW Systems," *Proc. ACM Symp. User Interface Software and Technology* (UIST 95), ACM Press, 1995, pp. 159-168.
3. W.C. Hill and J.D. Hollan, "Edit Wear and Read Wear," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 92), ACM Press, 1992, pp. 3-9.
4. C. Plaisant, "LifeLines: Visualizing Personal Histories," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 96), ACM Press, 1996, pp. 221-227.
5. W. Sack, "Conversation Map: A Content-Based Usenet Newsgroup Browser," *Proc. 5th Int'l Conf. Intelligent User Interfaces*, ACM Press, 2000, pp. 233-240.
6. M.A. Smith and A.T. Fiore, "Visualization Components for Persistent Conversations," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 2001), ACM Press, 2001, pp. 136-143.
7. S. Farnham, et al., "Personal Map: Automatically Modeling the User's Online Social Network," *Human–Computer Interaction* (INTERACT 03), IOS Press, 2003, pp. 567-574.
8. J.S. Donath, "A Semantic Approach to Visualizing Online Conversations," *Comm. ACM*, vol. 45, no. 4, 2002, pp. 45-49.
9. D. Kurlander, T. Skelly, and D. Salesin, "Comic Chat," *Proc. Siggraph 96*, ACM Press, 1996, pp. 225-236.
10. F.B. Viegas and J.S. Donath, "Chat Circles," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 99), ACM Press, 1999, pp. 9-16.

Social browsing tool requirements

Social browsing tools let users browse, become aware of, acquaint themselves with, and keep abreast of the interaction space's social milieu. Unlike the visual analytic tools that Web site administrators and designers use to monitor a site's usage and problems, the social browsers introduced in this article target a site's visitors and participants. We derived a number of user requirements for the design of social visualizations that support their visual analysis activities as part of their explicit social browsing activities. Some of these requirements are based on Donath's visual design and visual engagement principles[2] and Shneiderman's information-seeking mantra of overview first, zoom and filter, then details on demand.[3] Others were derived from empirical data that we collected based on the users' interest and usage of the CHIplace and Portkey Web sites.[4]

Legible visualizations

Visualizations are viable as user interfaces for casual users if they are intuitive, easy to use, easy to access, and visually engaging.[2] The key challenge is to identify visual design elements and visual techniques to depict the people, activity, and vitality in conversational spaces (see, for example, the Babble system[5]). Several systems address this challenge using simple, compact, and legible visualizations to highlight particular patterns and convey aspects of the social context (for example, tenor, activity, and purpose) in the online space. Both of our social visualization browsers extend these and other techniques to present aspects of a Web site's social context.

Encode interpretations of meaning

Explicitly incorporating interpretations of the data's meaning and relevance into the visualizations helps evoke an appropriate, intuitive interpretation from the users about the interaction context. Herein, the main challenge is to identify what data to represent and how to visualize it in an evocative fashion. ePlace exploits people's ability to interpret and navigate using maps in developing a novel, spatially organized, interactive Web site map that provides visibility of people, activities, and interactions.[6] Loom exploits people's knowledge about the social meaning of categories (such as provocateur or leader) and people's impression of different combinations of visual design elements to evoke a common interpretation of the data.[2] PeopleGarden uses a flower and garden metaphor that leverages people's ability to read the visual representation and to associate attributes of the metaphor—such as life and freshness conveyed by the position and color of flower petals—in interpreting the visualization (for example, multiple voices in or the democratic nature of an environment).[7]

Our People browser combines aspects of both the spatial organization and visual design approaches to help users interpret the visualization. As we'll discuss later, the eTree browser uses a different type of metaphor to facilitate an intuitive interpretation of its visualization content.

A community's public face

The key requirement of creating visualizations for participants of an online social space is to let them explore and understand the space's social environment. Frequently, a particular facet of the community is presented through its chat room activities such as cancer care support or young adult networking. PeopleGarden conveys such facets using visual representations of the amount of participation and the vitality in conversational spaces. Others, such as ePlace, try to convey the purpose and activities of the supported social context (such as customer-relationship management, auctions, and so on). In our work, we present visualizations that concentrate on different facets of group identity: a community membership map of CHIplace in the People browser and a conversational and participation portrait in the eTree browser.

Support change

Discussion sites grow organically as a result of new forums, threads, and messages. Similarly, social interaction sites grow in membership as new people join. Consequently, the design of social browsers must accommodate change. They also need to provide effective mechanisms to represent the salient information in the large volume of information that is created.

To accommodate evolution, we use a combination of metaphors, information presentation techniques, and user interaction mechanisms.[8] Specifically, we use dynamic queries and tool tips extensively to let users browse the information. The clustering technique, used in the People browser, is resilient to new data and can identify the key clusters of work roles. In the eTree browser, we exploit the botanical ecosystem metaphor to present the evolution of discussions. We also provide a feature that animates the evolution of tree growth over time.

Coupling with the browser

Users need techniques to deal with large volumes of information. Existing Web browsers are effective for providing details about community members and discussions. Social visualizations are useful for presenting an overview of an interaction context but their representations break down for presenting minutiae of detail. Our work illustrates how we can couple social visualizations with Web browsers to provide a social interface to the Web. The resulting system presents information about the social context and provides dynamic query functions to help users navigate the information in detail.

CHIplace People browser

We developed CHIplace, a social interaction Web site for the 2002 ACM Computer–Human Interaction conference, to support information exchange and interaction by the community of researchers and practitioners of human–computer interaction (HCI). The site operated from March 2001 until April 2002. In an earlier article, we presented the design, development, and evaluation of this social interaction Web site.[4] Here, we focus on a component of the site, the People browser (see Figure 2a), that we created in response to users' interest in learning about who the CHIplace members were.

Diversity and interest in members

From the outset, CHIplace provided tools for finding members of the site's community. For example, the main People page presented a list of people who joined most recently and a gallery of randomly selected pictures of members. This serendipitously introduced newcomers and members to each other. Users could also

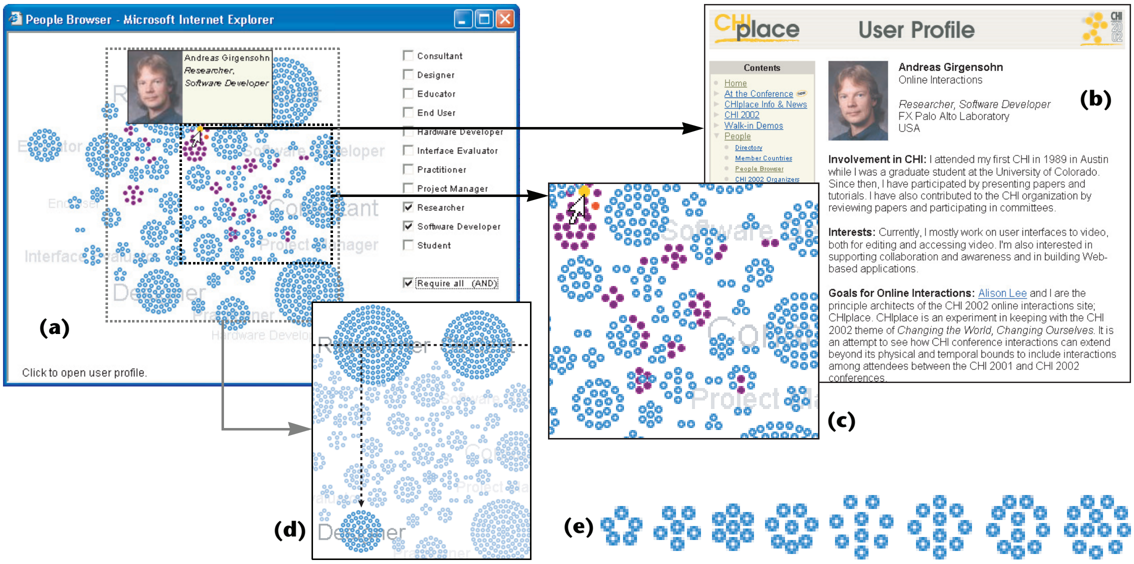

2 **(a) CHIplace People browser displaying about 1,200 users in a 400 × 400 pixel area—hovering over a user dot brings up a tool tip summary about the highlighted user. (b) A user profile Web page from the CHIplace site. (c) Enlargement of a portion of the People browser visualization revealing landmarks (role names displayed in different font sizes and levels of gray proportional to the number of their members) for roles. (d) The display is rotated to align the two largest clusters (excluding the cluster without roles) on a horizontal line and flipping the display so that the third-largest cluster is below the line. (e) Completely enclosed clusters for different sizes of clusters.**

select individuals to view the person's profile information (see Figure 2b).

We found, from the site usage log, that users often frequented the People section of the Web site. Some users kept reloading the main People page to see groups of random pictures. This suggested a need for a new functionality and prompted the idea of a visualization that would let visitors obtain a holistic view of the membership.

Among the many pieces of profile information that members provided, we found that the CHI roles profile field was the one optional field that most people filled out (92 percent). This field let members choose one or more of the 11 provided values or enter their own description of their involvement in HCI. In examining the selections that were made, we found that 47 percent of the members selected a single role while 45 percent of the people used multiple descriptors.

Given the near split between single and multiple roles selection, we could not present a visual community map based solely on grouping individuals in one of the 11 possible roles. Instead, this suggested a visualization that identified the natural clusters existing among the groups of CHI roles selected by CHIplace members. While the data is one-dimensional (values of the selected roles), the fact that users can pick several values makes it impossible to place users along a single axis. From a usability and ease of interpretation perspective, we felt that a 2D map was accessible to our users. Because this map only depends on distances between users, additional attributes (such as country of residence) could be included in the visualization provided that those attributes could be incorporated into a single distance measure that expresses differences between two users as a numeric value.

Visual community map

We produced a simple and legible community map by combining automatic clustering and 2D layout techniques as well as visual design elements and display enhancement techniques. We grouped members by the sets of roles that they have in common using hierarchical agglomerative clustering. The people clusters are displayed as concentric rings of dots in a 2D space so that related clusters are closer together than unrelated ones. Each dot represents a CHIplace participant. We used a graph layout algorithm based on a spring model to place nodes representing clusters of people (see Figure 2a). As new members join and current members update their CHI roles, the visualization layout accommodates changes in clusters, membership, and size.

We used numerous display enhancements to facilitate the users' interpretation and comprehension of the visualization. First, we determined empirically that no more than 80 clusters should be displayed for ease of understanding because displaying more clusters led to a scatter-plot-like visualization without clearly separated clusters. Next, we added landmarks behind the corresponding clusters. This lets users associate the roles for the clusters (see Figure 2c). Then, we rotated and flipped the display so that the three largest clusters are in a consistent configuration (see Figure 2d). This provided a reference guide for exploring the visualization in detail and as a whole. Finally, the dot of the logged-in user appears in red (see Figure 2c) to let them see where they are on the People map.

While the current display could scale up to more users, a very large user population would require enhancements to the visualization techniques. For example, instead of displaying an individual dot for each user, clusters could be visualized as circles with areas proportional to the number of cluster members. This kind of visualization would also allow for the display of more clusters because clusters could overlap partially without losing their visual separation. Clicking on one of those circles could zoom the user in on the details of the cluster through selective enlargement of the cluster or display of a list of the members' names. In the latter case, the list could either contain additional information about the users or the current tool tip approach could provide that information. Featuring an overview first that lets users zoom in and filter content might reduce serendipitous discovery because of reduced detail but it would make it easier to explore clusters with many members.

We also used several strategies to make the visualization appealing and to display a large number of people in a small display footprint. Our first strategy was to specify a minimum distance between dots to provide sufficient visual separation while keeping the clusters as tight as possible. A second strategy was to display the clusters as closed and complete structures to enforce visual unity of the group of dots for the cluster. We accomplished this by determining for every cluster, the number of dots in the innermost ring of the cluster that are spaced by the minimal dot distance. This approach results in a completely filled outer ring (see Figure 2e).

Finally, the visualization supports numerous functions that facilitate exploration. First, when the cursor hovers over a user dot (see Figure 2a), the dot changes to a solid yellow and a tool tip pops up showing the picture, name, and CHI roles for the highlighted individual. Second, the overview visualization supports browsing of the CHIplace member profiles. Clicking on a dot opens the CHIplace user profile Web page for the person (see Figure 2b). Viewed user dots appear in a lighter color shade similar to visited links on a Web page. Finally, checking the roles in the checklist will highlight, in dark purple, the members with some or all of the roles selected in the checked boxes (see Figure 2c for a closer view of Figure 2a).

Implementation

We implemented the People browser as a Java applet. To quickly transfer people's pictures to the applet, we created Zip archives containing the pictures for groups of 100 people. Depending on the number of people with pictures in a group, the sizes of these archives can range from 30 to 150 Kbytes. The first time a picture from a previously not shown group is displayed, the applet loads the corresponding archive and caches the compressed images. Images can be uncompressed on the fly for display without any noticeable delay.

The clustering and spring model simulation steps—discussed next—are performed periodically on the server. One reason is that the computations take several seconds for about 1,200 people. Another reason is that the data changes less frequently compared to visualization requests. The applet loads the precomputed data when it starts.

We use the Jaccard coefficient as the similarity measure in the clustering step. For the sets of roles R_i and R_j for two people, we compute the Jaccard coefficient C_J as follows:

$$C_J = \frac{|R_i \cap R_j|}{|R_i \cup R_j|}$$

People without roles have the maximum distance to anyone with roles.

A node in a graph represents each people cluster. We use a spring model to determine an optimal layout in which the spring system is in a state of minimal energy. We chose this model for three reasons. First, the derivatives guide the simulation in a direction that converges quickly to a satisfying state. Second, because only distances are needed, we can use a non-Euclidian space. Third, unlike self-organizing maps,[9] we can place nodes in any position and not just in a grid or similar structure of cells.

Our spring model assumes a fully connected spring system—that is, a spring between every pair of nodes. The length of the springs at rest corresponds to the desired distance between the nodes. To avoid finding only a local minimum, we start the spring model in different random states and refine the state with the lowest energy. We used the same similarity measure from the clustering step for the desired distance between two nodes. To ensure that large clusters are placed more accurately than the small clusters, we set the spring constants to be proportional to the number of people in the clusters. We used the same random number seed to initiate the spring model simulation to ensure that the same layout is produced for the same data set rather than a random layout variation.

Comparison to Visual Who

The People browser is similar to the Visual Who system,[10] which provides a graphical browser for interactively examining the social groupings of community members. Like the People browser, it uses visual design techniques to create a legible visualization.[2] Visual Who does not use clustering to determine groupings but lets users place anchors representing groups of interest. The system then uses the spring simulation to generate a graph layout by animating group members' names gravitating toward the corresponding anchors. This results in a simpler spring simulation because springs only connect movable nodes to unmovable anchors. Also, Visual Who uses a different similarity measure of people that is based on shared group membership.

The People browser uses a dynamic query interaction style that encourages familiarization through a peek and poke interaction model. Visual Who encourages discovery through active user exploration by selecting the categories of interest to be visualized. Thus, Visual Who requires greater user experimentation and visual sophistication compared to the People browser.

Finally, the People browser provides all the clusters in one view, in a smaller display footprint, and uses visual feedback and query techniques to supplement the overview with details. In contrast, Visual Who's style of

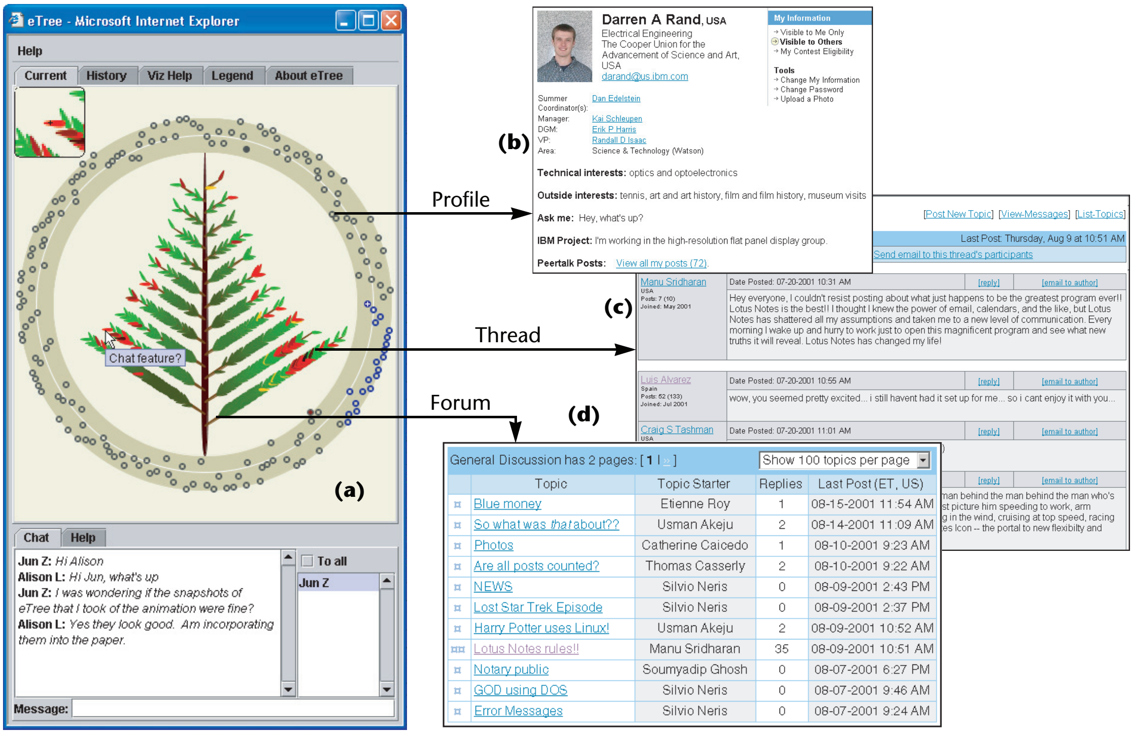

3 (a) eTree browser. (b) Each user circle around the larger ring corresponds to a user profile Web page on Portkey. (c) Each leaf corresponds to a discussion thread page on Portkey. (d) Each tree branch corresponds to a discussion forum on Portkey.

exploration of selected groups lacks an overview mechanism. It requires a larger display area to show people's names instead of dots and even with the larger display, the names can obscure each other.

Portkey eTree browser

We created Portkey, a social interaction Web site, to support the summer 2001 interns in their transition to and their work assignments at IBM T.J. Watson Research Center. We previously discussed the site's design principles, design, social capabilities, and usage patterns.[4] Here, we present one site feature, the eTree browser, which we created in response to the interns' interest in getting to know one another through the discussion forums.

People and discussions

We had two design goals in developing Portkey. The first was to create an online environment that would support the sharing of relevant personal experiences among interns. We provided various forms of asynchronous discussion forums. By the end of the summer, 156 students wrote 1,677 messages with an additional 73 messages contributed by 19 IBM staff members. The posters grouped those 1,750 messages into 840 threads.

A second Portkey design goal was to facilitate the development of the personal and professional social networks that the interns needed to function effectively in their new environment. We developed mechanisms that made the members visible to and identifiable by others such as linking posts with the poster's profile page and providing a people directory (categorized, for example, by university affiliation, first name, and technical disciplines) and a people gallery similar to CHIplace. An examination of the site usage log revealed that about 80 percent of the page accesses were to the discussion and people pages.[4]

The interest in people and discussions prompted us to develop the eTree browser to address the lack of a holistic view of or browsing tool for the social context.

Ecosystem metaphor

We used an ecosystem metaphor—consisting of the different parts of a tree—to create the eTree visualization. We used this metaphor to map the discussion structure onto various parts of a tree such as the trunk, branches, and leaves (see Figure 3a). We mapped threads in a forum to leaves. Hot topics are highlighted as yellow leaves. Posters are represented as colored circles (gray for interns and blue for IBM staff members) that are placed around the perimeter of a ring centered on the tree. A poster's position in from the border of the ring is the square root of the number of posts. The angle of the circle's position is determined by a more arbitrary measure, the user ID.

The eTree in Figure 3a visualizes 175 users. While that number could easily be increased up to a factor of three in the same space, a much larger number of users could be displayed in clusters as discussed for the People browser. Furthermore, if it is desirable to position posters who participate in the same discussion threads close to each other, we could determine the distance between two users by the number of discussion threads in which they participated.

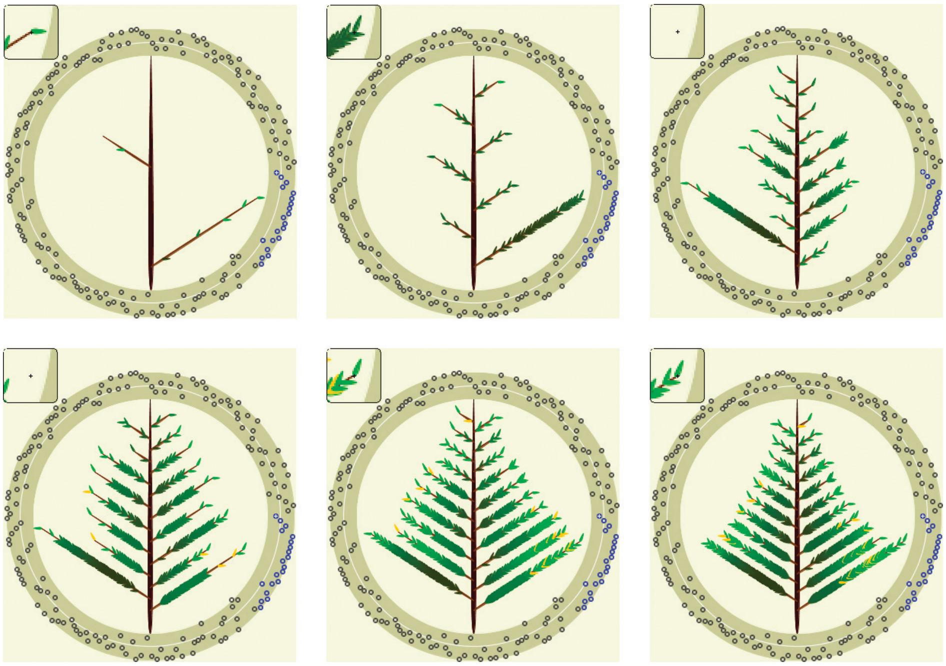

4 **eTree animation of discussion forums**

eTree's visualization evolves organically in response to new posters, posts, threads, and discussion forums. New posts appear as new leaves in light green. The leaves of older posts turn dark over time. Figure 3a shows Portkey's eTree with 21 branches, one for each of the 21 discussion forums. The tree's shape, branch length, and leaf density illustrate the number of forums and the popularity of certain forums and threads. It also reveals that most of the Portkey posters were interns (gray circles) with only a handful of IBM authors (blue circles).

In its current form, the visualization does not scale up to hundreds of discussion forums containing hundreds of threads. The tree does not have room for many more branches. While the magnifying glass feature helps with dense leaves, threads cannot be navigated once the distance between leaves is less than one pixel. To deal with this problem, we envision two sliders that function like scroll bars. The first slider would control the time range. Rather than displaying all threads in a discussion forum, only those within a certain date range would be shown. Moving the slider would move the thread leaves up and down the branches with leaves appearing and disappearing at the ends. To avoid disorientation, branches should not change position if the number of their visible leaves changes. The other slider would flip through discussion forums and map different forums to the tree branches.

The eTree browser incorporates several awareness cues. Newcomers can learn about the evolution of the discussions from an animation of the tree's growth (see Figure 4). Users can also see other users copresent on the site as filled-in circles around the ring's perimeter. The location of users in Portkey's discussion area is indicated as an additional filled-in circle on the tree trunk, branch, or leaf, which corresponds to the forum listings page, discussion forum, or thread, respectively. If the user is also a poster, his or her filled-in circle has a crosshair in it (see the topmost blue user circle around the ring in Figure 3a). These awareness cues let co-present visitors start a text chat with each other.

As with the CHIplace People browser, Portkey eTree browser users can employ tool tips to obtain additional information about an element (see the "chat feature" tool tip in Figure 3a). Also, clicking on the visual elements brings up the corresponding Web page that provides details (see Figures 3b, 3c, and 3d). Users can also query a visual element (see the sequence of user queries in Figure 5). These queries can also link to Web site pages (see Figure 6). The eTree browser includes a magnifying lens anchored in the upper left-hand corner (see Figure 3a). This lets users pick out individual threads (represented as a leaf) from the large numbers of threads.

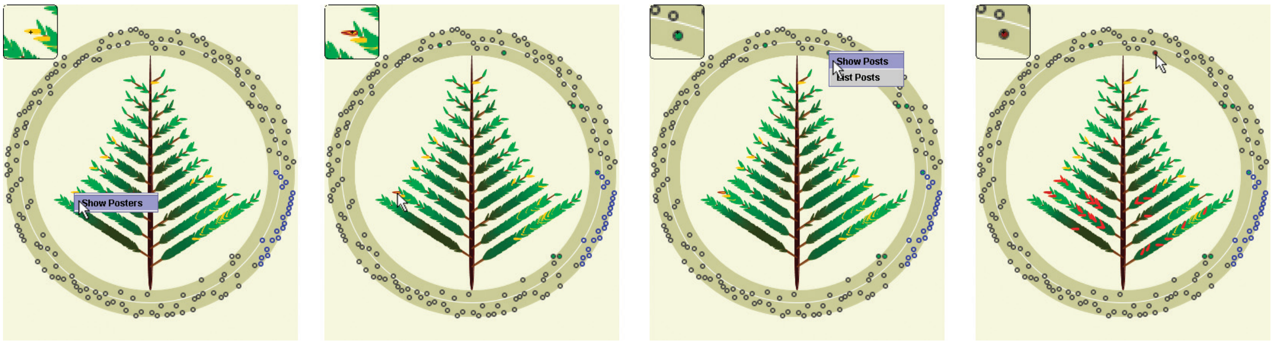

(a) **(b)** **(c)** **(d)**

5 **(a) Right-clicking over a leaf brings up a menu to request all the posters of a thread. (b) Query results for all posters of a thread. (c) Right-clicking over a given poster brings up a menu to request all the posts by a poster. (d) Query results for all the posts by a given poster.**

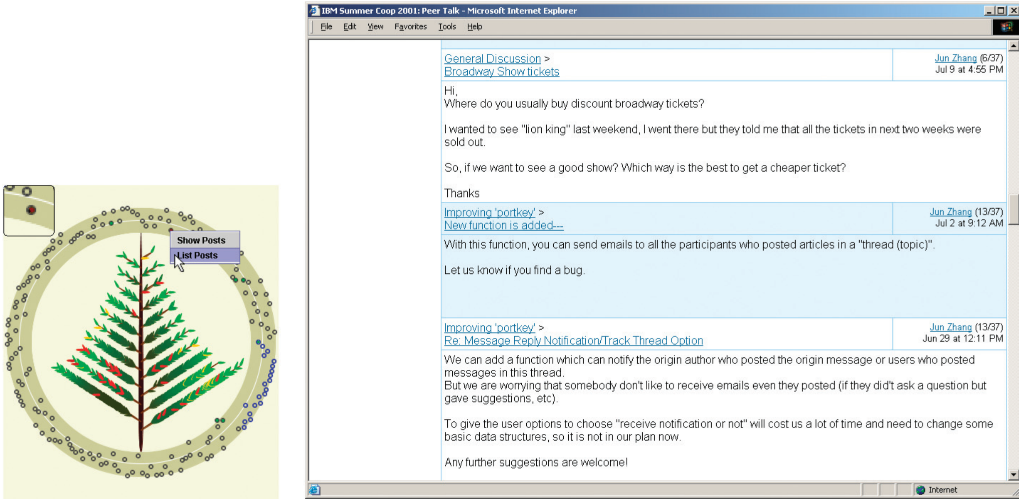

(a) **(b)**

6 (a) Requesting all the messages posted by a given user results in navigation to (b) a page on the Web site containing a list of messages posted by the given user.

Implementation

We implemented the eTree browser as a Java applet. The applet connects to a server and exchanges information via Java object serialization. The server periodically retrieves updated information about discussion forums, registered users, and site visitors and distributes that information to all connected eTree browsers. For performance reasons, the server checks less frequently for information that is less likely to change such as posts in discussion forums.

The eTree browser exploits several strategies to use the available display space and to highlight important information. Branches of the tree representing discussion forums are spaced evenly along the tree's trunk. The number of threads in the corresponding discussion forum determines the branches' length. The discussion forums are sorted by the number of threads so that the longest branches are at the bottom of the tree. Using the square root of the number of threads for the branches' relative length produced visually more appealing results than making the length proportional to the number of threads, because short branches could be avoided.

Leaves representing the threads in a forum are placed along the branch so that newer threads, in light green, lie at the branch's outer edge. We initially considered using a fish-eye distance for the spacing between leaves so that there would be more space between the newer, and presumably more interesting, leaves. However, we found that this approach made it too difficult for users to select leaves of the older discussion threads. We also found that spacing leaves solely by their temporal distance produced large gaps between some leaves and caused tight spacing between others. Instead, we combined the two approaches and used the distance between two leaves as the average of a temporal and a constant distance.

To visualize the age of discussion threads, we used different shades of green. The green component of the RGB color is varied between 25 and 100 percent with the age of threads so that an older thread's color is darker than that of a newer thread. Leaf color also contains a 20 percent red component so that the color of older threads moves toward a shade of brown. Bright yellow leaves indicate active threads with many messages. To visualize all posts by a selected user, the corresponding threads are red. Threads that have many other authors are shown in a darker shade of red to make the hot-topic threads more visible.

Comparison to PeopleGarden

eTree's botanical metaphor is reminiscent of the flower and garden metaphor used in PeopleGarden.[7] Both metaphors provide life-like portraits of an online social environment. PeopleGarden portrays each participant's contribution in a newsgroup as a petal in the person's flower with all participants in the newsgroup grouped in a garden.

The two systems have four main differences. First, PeopleGarden provides a participant-centric view while the eTree browser provides both a view of posting content as the tree and of participant involvement represented by circles around the tree. In addition, eTree's social environment includes all the discussion forums and active participants while each PeopleGarden visualization represents a single newsgroup, chat room, or discussion forum. eTree then has the benefit of providing a view of the larger context and allows scaling of the visualization to accommodate new discussion forums and threads. Second, PeopleGarden provides a snapshot of the environment while the eTree browser is dynamic and evolves as the social environment grows. Third, the eTree browser provides awareness of peers' online status (a solid circle) while PeopleGarden only reflects post history. Finally, the eTree browser provides many query and exploration tools to drill into the details including the linking of visual elements to Web pages on Portkey.

Discussion

In creating and deploying the two browsers, we learned several lessons. These lessons suggest areas for further improvement and valuable insights for others developing similar social browsers.

User feedback and browser usage

We introduced the Portkey eTree browser shortly before the end of the interns' stay so we only have anecdotal evidence of peoples' feelings about it and informal feedback about the metaphor. Many people found the visualization attractive; in particular, they liked the ability to view the growth and evolution of the discussions. However, a main concern was privacy, which we'll discuss later.

For the CHIplace People browser, we collected extensive usage data.[4] In the interest of privacy, and to honor the site's privacy disclosure that we only track site usage, we did not program the People browser applet to track how people used the applet features. However, the site usage log includes information about sessions where users employed the People browser and when they clicked on circles to find out more about people.

The People browser was available during the last 23 of the 57 weeks of the CHIplace's operation. During that period, it was used in 9 percent of the 17,126 sessions and by 11 percent of the 10,600 IP addresses. That usage is comparable to usage of the main people page showing the picture gallery in 14 percent of the sessions. When only considering the 4,202 sessions with above average page accesses (8), the People browser was visited in 25 percent and the main people page in 41 percent of those sessions. While the People browser did not replace the previously available tools for browsing people, it was a useful addition.

The announcement of the People browser caused an overall usage spike (55 percent more weekly sessions than the average of the four weeks before). During the introduction week, 35 percent of the sessions used the People browser. For the following 10 weeks, the browser usage ranged from 9 and 27 percent of the weekly sessions. Usage leveled out between 6 and 8 percent of the weekly sessions in the last 12 weeks.

Roughly half the sessions in the observed time were from repeat visitors—that is, addresses that appear in more than one session (8,989 sessions and 2,063 addresses). Among repeat visitors, 11 percent of the sessions and 28 percent of the addresses used the People browser. In the sessions that used the People browser, 9 percent clicked on a person's circle to get more information (on average 6 clicks in those sessions). In 36 percent of the sessions that used the People browser, users opened the browser in a pop-up window. Those users (16 percent) were more active in clicking on person dots (8 clicks).

Privacy

Privacy is a key concern with providing social browsers to highlight patterns within a social space. The information portrayed by both the People and eTree browsers was based only on public information in the CHIplace people profile and Portkey discussion forums. The information was either volunteered or drawn from traces of people's activities in the social space. In the latter case, users were aware that the discussion forums were part of Portkey's community data and that the shared information was visible to members of the social space. This policy was consistent with other informal content created and shared by Portkey members.

Of course, in making some of the information more salient in the browsers, the browsers call attention to information that was buried with other data; thus, participants might view the highlighted data as being sensitive information. Therefore, to avoid the potential of a misunderstanding of the public and private nature of information that people share, it's important that the social space provides cues to its participants about the boundaries of the community's public spaces. This provides people with cues about when they are in public or private spaces so that they can act accordingly. We did not receive complaints or explicit feedback from users that would suggest any problems with privacy issues.

Facets of group identity

A consequence of developing visualizations that evolve through the interactions of participants is that it allows the collective contributions to define the emergent character of the community shaped in part by the visual representation. In the case of the People browser, the size and distribution of the concentric clusters of CHI roles existing in CHIplace provides an abstract representation of the CHI roles. The shape of the life-like tree, number of contributors, and number of posters closer to the circle's center characterize Portkey's social vitality. In different online social spaces, these visualizations would likely yield different visual results. Hence, they present interesting portraits of their communities.

User control

Coupling the social visualization with visual feedback and query mechanisms provides two benefits. First, the social milieu is visible to visitors for exploration. Second, participants of the social space can also explore the visualizations to see how they are seen by others and how their contributions influence the community portrait's character. The reciprocity permits members to make adjustments to what they reveal about themselves. For example, CHIplace users could make refinements to their self-descriptions of CHI roles and thus control the presentation of their electronic self on CHIplace. Portkey users could use the reflection to examine how they might be viewed by others based on the level of their participation and their interactions with other members.

Technical considerations

The social information in the visualizations were processed for layout and presentation and translated into a format suitable for communication with an applet. For the eTree browser, it was inefficient to extract, in real time, the information about visitors' location from server logs. Instead, we programmed the header portion of each Portkey page to keep track of visitors. Because of the ephemeral nature of visitor locations, we kept the information in memory and removed it after a few minutes of visitor inactivity. We believe that similar approaches are appropriate for tracking visitors on many sites with dynamically generated pages.

Usability considerations

We received user feedback on the usability of both browsers. The People browser's layout algorithm sometimes places the nodes in nonintuitive locations. The algorithm does this because the global optimization technique of the spring simulation finds a configuration in which all springs connecting pairs of clusters reach their desired lengths. So, sometimes clusters might get pushed significantly out of the expected position. We are considering using fewer springs and fairly localized repulsive forces. The People browser users also requested a visualization that slowly changes with the addition of new users. Given that the current layout algorithm is sensitive to small changes, we are investigating a different algorithm.

While both browsers offer tool tips to identify items in the visualization, it would be useful to have more landmarks. The eTree browser could adapt the People browser's approach of labeling the clusters of the major roles by labeling the tree branches. Both browsers would benefit from a search facility to find people and messages.

Conclusions

In addition to the novel visualizations and representation of two facets of a group's identity, our work has three other notable contributions. First, the visualizations organically evolved as a consequence of people's contributions. Second, we used two different approaches to create the visualization: People browser combines automatic techniques with visual design elements and information visualization techniques and eTree uses an ecosystem metaphor. Finally, each social browser was integrated with other functionality at the social site to let users obtain overview and detail, see how others view them, and provide them with information that they can use to guide their participation and possibly regulate their behavior.

User feedback and comments provide strong support of the value of the two browsers. We will need further studies to examine how such techniques can alter people's participation and behavior in Web communities. We believe that such browsers enhance the user interface to and the social interface from Web communities. They are also novel technologies that differentiate a social Web from the traditional Web; particularly when the Web is used as a social technology. ■

References

1. R. Oldenburg, *The Great Good Place: Cafes, Coffee Shops, Community Centers, Beauty Parlors, General Stores, Bars, Hair Salons and Other Hangouts at the Heart of a Community*, Marlowe and Co., 1989.
2. J.S. Donath, "A Semantic Approach to Visualizing Online Conversations," *Comm. ACM*, vol. 45, no. 4, 2002, pp. 45-49.
3. B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations," *Proc. IEEE Symp. Visual Languages*, IEEE CS Press, 1996, pp. 336-343.
4. A. Girgensohn and A. Lee, "Making Web Sites be Places for Social Interaction," *Proc. ACM 2002 Conf. Computer Supported Cooperative Work*, ACM Press, 2002, pp. 136-145.
5. T. Erickson, et al., "Social Translucence: Designing Social Infrastructures that Make Collective Activity Visible," *Comm. ACM*, vol. 45, no. 4, 2002, pp. 40-44.
6. Y. Jung and A. Lee, "Design of a Social Interaction Environment for Electronic Marketplaces," *Proc. Conf. Designing Interactive Systems* (DIS 2000), ACM Press, 2000, pp. 129-136.
7. R. Xiong and J. Donath, "PeopleGarden: Creating Data Portraits for Users," *Proc. 12th Ann. ACM Symp. User Interface Software and Technology*, ACM Press, 1999, pp. 37-44.
8. C. Ahlberg and B. Shneiderman, "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 94), ACM Press, 1994, pp. 450-456.
9. T. Kohonen, *Self-Organizing Maps*, Springer Verlag, 2001.
10. J.S. Donath, "Visual Who: Animating the Affinities and Activities of an Electronic Community," *Proc. 3rd ACM Int'l Conf. Multimedia*, ACM Press, 1995, pp. 99-107.

Alison Lee *is a research staff member at the IBM T.J. Watson Research Center. Her research interests include tools, applications, user interaction mechanisms, and software architectures that enrich human–computer interaction, human–human interaction, and end-user accessibility experiences. Lee has a PhD in computer science from the University of Toronto.*

Andreas Girgensohn *is a senior research scientist in the Interactive Media group at the FX Palo Alto Laboratory. His research interests include design, development, deployment, and evaluation of applications for video editing and indexing and for organizing digital photos. Girgensohn has a PhD in computer science from the University of Colorado at Boulder.*

Jun Zhang *is working toward a PhD in human–computer interaction at the University of Michigan. His research interests include social network-based information seeking, visualization of online communities, and how information technology (such as email technology) affects productivity in organizations. Zhang has a BS in engineering science from Tsinghua University, Beijing, and an MS in information science from the University of Michigan.*

Readers may contact Andreas Girgensohn at FX Palo Alto Laboratory, 3400 Hillview Ave., Palo Alto, CA 94304; andreasg@fxpal.com.