

IBM Research Report

Socially Augmenting Employee Profiles with People-Tagging

Stephen Farrell, Tessa Lau, Eric Wilcox, Stefan Nusser

IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099

Michael Muller

IBM Research Division
One Rogers Street
Cambridge, MA 02142



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Socially Augmenting Employee Profiles with People-Tagging

*Stephen Farrell, Tessa Lau, Eric Wilcox,
Stefan Nusser*
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120 USA
+1-408-927-1856
sfarrell@almaden.ibm.com

Michael Muller
IBM Research
One Rogers Street
Cambridge, MA 02142 USA
+1-617-693-4235

ABSTRACT

Employee directories play a valuable role in helping people find others to collaborate with, solve a problem, or provide needed expertise. Serving this role successfully requires accurate and up-to-date user profiles, yet few users take the time to maintain them. In this paper, we present a system that enables users to *tag* other users with key words that are displayed on their profiles. We discuss how people-tagging is a form of social bookmarking that enables people to organize their contacts into groups, annotate them with terms supporting future recall, and search for people by topic area. In addition, we show that people-tagging has a valuable side benefit: it enables the community to collectively maintain each others' interest and expertise profiles. Our user studies suggest that people tag other people as a form of contact management and that the tags they have been given are accurate descriptions of their interests and expertise. Moreover, none of the people interviewed reported offensive or inappropriate tags. Based on our results, we believe that people tagging will become an important tool for relationship management in an organization.

ACM Classification H5.3 [Information interfaces and presentation]: Group and Organization Interfaces - CSCW.

General Terms Design, Human Factors

INTRODUCTION

Enterprise directories are an important tool for finding information about other people in the workplace to build awareness and to learn about them prior to contacting them. Our corporate directory, known as "BluePages," fields more than 1.5 million hits per day and is widely regarded as the most successful Intranet application.

Users come to the directory not only for contact information but also to learn more about the people they work with. In support of this behavior, BluePages has an elaborate profile

template that allows users to specify their projects, expertise, experience, teams, education, and other attributes. Despite the overall success of BluePages, the level of individual contribution to profiles is frequently cited as a problem. Current methods for populating and maintaining directory content have many limitations. Critical parts of the profile, built automatically from human resources databases (e.g., the organization chart), are reliable and up-to-date. Other parts such as a biography, experience, teams, projects, and interests rely on individuals filling out that information for themselves and are less reliable. In the BluePages directory, only 40% of profiles have been updated in the past nine months and 22% have been updated in the past three months. These numbers reflect *any* changes to the profile; the number of employees who have provided comprehensive and up-to-date information is far lower.

We realized that the problem of finding information about people in an organization is conceptually similar to finding web pages on the Internet. Could we take ideas from social bookmarking/tagging systems like del.icio.us [5] and Dogear [8], in which millions of users have tagged web sites with descriptive words, creating a collectively-authored folksonomy of content? If we provided a system in which people could tag *people*, could we use these tags to create a folksonomy of employees?

Social bookmarking systems provide a clear incentive for users to participate. Typically bookmarks are centrally managed by the service provider without a fee to the user, accessible from any web browser on any connected workstation, and searchable for later recall. People engage in social bookmarking not only for altruistic reasons but because of these concrete benefits [7]. However, as a side-effect of this behavior, they create meaningful tags and rankings for the underlying resources, providing labels that other people can use to find that information. These tags often tend to be both more current and more complete than automatically generated or owner supplied classifications. What makes these systems remarkable is their ability to leverage the self-serving actions of individual users to provide a significant benefit to the user community as a whole.

We designed the people-tagging functionality to support contact management through bookmarking people. Beyond the

self-serving benefits, however, we believe that tags can be leveraged to augment employees' original profiles with additional information contributed by the community. Instead of relying on each employee to take the time to keep his own profile up to date, social tagging may leverage the work of a few active taggers in describing the interests and expertise of many others.

The objective of the work described in this paper is therefore to leverage the community benefits of social bookmarking in order to create accurate profile information. We have added a people-tagging feature to our enhanced employee directory, Fringe, which enables people to tag other people with key words describing their expertise or interests. The tagging feature has been available for over 13 months, with 3,462 users tagging 29,484 users with 73,438 tags to date.

While social tagging has been successfully applied to objects such as web pages, images and places, tagging other users—tagging *people*—presents new challenges. While objectionable tags can be an issue with other resources, users are particularly sensitive to information that appears in their profile page, especially in a professional context. It has been suggested that an approval process or way to remove unwanted tags would be required. We hypothesized that an open and transparent design would encourage good behavior. Furthermore, we designed the system so that all tags are traceable to the professional identity of the tagger, making users accountable not just to the community but to their employer. Will fears about inappropriate tagging become a barrier to the adoption of our system, or will users respond to unobtrusive social safeguards and tag unto others as they'd have others tag unto them?

After a survey of prior work, we begin with a description of the people-tagging capability that we have integrated into our next-generation employee directory, Fringe, as well as integration points with other applications such as email and instant messaging. We present some statistics characterizing how the use of tagging has increased over time. We then present a statistical analysis of tags as contributing novel information to existing profile data. Next, we report the results of a survey of 63 users to gather broad insight on how people perceive the people-tagging feature, as well as in-depth interviews with 19 users to understand specifically how people are using tags. Our results suggest that the people-tagging feature is useful for contact management, supports collective maintenance of the employee directory, and appears to have built-in safeguards against undesired tags. Finally, we conclude with our directions for future work.

PRIOR WORK

Social bookmarking systems have emerged recently as popular tools for organizing and sharing information. The del.icio.us social bookmark manager is one of the most successful of these [5]. Dogear illustrates how such a system can be adapted to the enterprise [8]. Since then, other systems have incorporated tagging as a means for annotating information, such as photos¹ and blogs².

¹<http://flickr.com>

²<http://technorati.com>

The concept of tagging *people* as opposed to web resources has received relatively little attention. Tagalag³ is one such system; it enables users to tag others based on their email address. Tagalag provides integration with web-based mail systems via a Greasemonkey script that adds tags to the web-mail interface. However, it seems to have had relatively little adoption. 43people⁴ is another people-tagging service that focuses on tagging celebrities. Some social networking sites like Xing⁵ enable people to create private tags, but not share them with others. To our knowledge, ours is the first social people-tagging system in widespread use.

One of the potential benefits of tagging is locating people in an enterprise based on their interests or expertise. Some expertise-finding systems have examined the tags individuals apply to social bookmarks as a measure of one's expertise. For example, John and Seligmann describe the ExpertRank algorithm for finding expertise based on bookmarks one has tagged [6] (see also [13]). In contrast, the tag rank analysis we have presented in this paper is based on tags applied to individuals *by their peers*. Another class of expertise-finding systems mines user behavior (e.g., code check-ins) to infer expertise [9]. In contrast, our system relies on the contributions from the community to develop associations between people and topic areas.

One use of people tagging is to help maintain relationships with one's contacts. ContactMap [11], based on a study that found that people have difficulty "remembering the identities of people in their social networks, particularly those who are important but contacted infrequently," addressed this issue with a spatial map. This representation is not easily aggregated for community value.

A SYSTEM FOR PEOPLE-TAGGING

Contact management is important for many people in professional and other settings. People want to keep track of others in order to find them again when needed, to remember details about them, and to organize them into useful categories. Just as social bookmarking is enabling people to keep track of webpages they encounter in the process of their work, we imagine people-tagging will enable people to keep track of people the same way.

To this end, we have developed a people-tagging service that has been integrated into our next-generation employee directory, Fringe. Moreover, it is extensible so it can be easily integrated with the many other applications where people encounter each other, such as email and instant messaging.

Design principles

We followed several design principles when designing the people tagging interface.

The first significant design decision was to allow users to tag one another without the express permission of the person being tagged. Many popular social networking services such as

³<http://tagalag.com>

⁴<http://43people.com>

⁵<http://xing.com>

MySpace⁶ and LinkedIn⁷ require that users send and accept invitations to establish connections. While Fringe also supports these invited connections, we wanted the tagging feature to be more lightweight and require no more effort than it does in applications like Flickr⁸ or Del.icio.us⁹.

We wanted our system to support constructive tagging practices without the need for policing or heavy-weight opt-in processes. In our design we leveraged understood principles that affect behavior. For example, in keeping with Goffman's concept of enacting a professional performance [4], people tend to avoid compromising their reputation with unprofessional meanness or derogation. Therefore, we have designed our system so all actions in our system are traceable to the professional identity of the actor. Users must authenticate with the corporate directory in order to create or modify any data. In support of "social translucence" [1], the authors of a tag are displayed when a hovering over the tag in the incoming tag cloud. Moreover, there is an outgoing tag cloud on each user's profile, so one cannot put a term on someone else's profile without simultaneously adding it to his own.

Beyond blatant misuse, we also wanted our systems to encourage users to take unintended or subtle consequences of their tagging behavior into account. For example, tagging someone "cobol-guru" might result in that person being asked a lot of questions about a programming language he might no longer use, even though the tag is accurate. A unique quality of people-tagging in contrast to other tagging applications is that when you tag someone, he can tag you back. In designing such a system, we expected users to take extra consideration and think about how they would want to be tagged before tagging others.

We also wanted to minimize the steps required for people to tag one another. One of the pain points in the existing corporate profiling system has been the amount of time and complexity needed to update personal content. With tags, we aimed to find a lighter-weight way to let people contribute content. We set out to make tagging as simple as entering a handful of keywords into a text field. Users can immediately see the impact of their contribution.

Fringe profiles

Each profile page in Fringe follows a standard layout as shown in Figure 1. The main center column on the page displays contact information for the profiled person at the top. Subsequent sections display automatically-gathered information about the person, such as the communities to which they belong, some key fields from their BluePages profile, recent blog entries and bookmarked web pages. The right column shows relationships to other people through invited connections and the organizational chart.

The tagging functionality is integrated into this directory interface. At the top of the left-most column is a small panel that provides an entry field for tagging this person (A). Tags are suggested in a drop-down box as the user types based

on overall tag frequency. All the tags ascribed to the profile by the current user are shown in green above the entry field. Tags given to this person by the current user may be discarded by clicking the "x" to the right of the tag. All tags are stored on a central server and persist across sessions. Subsequent visits to this person's profile page will load any tags attributed to the profile by the current user and display them in this panel. There is no way to create hidden or private tags.

Each tagger may ascribe multiple different tags to the same taggee. Tagging the same person with the same term twice has no additional effect. Multiple people may tag the same person with the same term; in this case, that tag is given more visual weight in the tag cloud view.

Three *tag clouds* are displayed beneath the tagging widget. The first tag cloud ("tagged by -- people") represents a collection of all the tags given to the profiled person, by any user (the "incoming" tag cloud). As is the convention, tags are sorted alphabetically and typographic weight and size are used to signify frequency, with tags that were given by multiple people appearing in larger and bolder type. Hovering the mouse over any given tag provides a count and the list of users who ascribed the tag (B). The second tag cloud follows the same interface conventions as the first. However, the contents of this cloud represent the tags that the profiled person has ascribed to others (the "outgoing" tag cloud). In each case, a count of the number of people (number of people tagging in the first cloud, number of people tagged in the second) is called out at the top of the panel as a hyperlink. Clicking on either of these links will display the list of users who have tagged this person, or whom this person has tagged, respectively. A third cloud displays a tag cloud of the tags this person has used in our colleagues' social bookmarking system, Dogear [8].

Users can pivot on tags by clicking on them. Pivoting on a tag results in a display of all the people across the company tagged with the selected value (Figure 2) sorted by the frequency with which they have been ascribed that tag. The same results may be reached by typing the tag text into the top-right corner search field. A tag cloud for all the people in this list is shown in the left column, which may be used for further pivot-browsing. Additionally, the email addresses for these people are displayed below the tag cloud, in a form that can be copied and pasted into an email. A checkbox on the left side of the page allows the user to filter the list to only those whom she has tagged, or to show all users with the given tag. The default view is the "biz-card" view, which shows a list of photographs of each of the people in the group. The geography view displays a mashup of a Google map with the addresses of people in the group, enabling one to see where all the members of the group are located. The network view displays a social network diagram based on tagging and friending data.

Integration with other applications

References to people appear everywhere: in email, on the web, in instant messaging, in documents. To make people-tagging as easy as possible, we wanted to provide the ability to tag people *in context*, without the need to launch a separate web browser to access Fringe in order to tag them.

⁶<http://myspace.com>

⁷<http://linkedin.com>

⁸<http://flickr.com>

⁹<http://del.icio.us>




Figure 1: Fringe profile page. (A) Tag widget with auto-suggestion. (B) Tag clouds showing tags this person has used on others, and tags used by other people on this person.

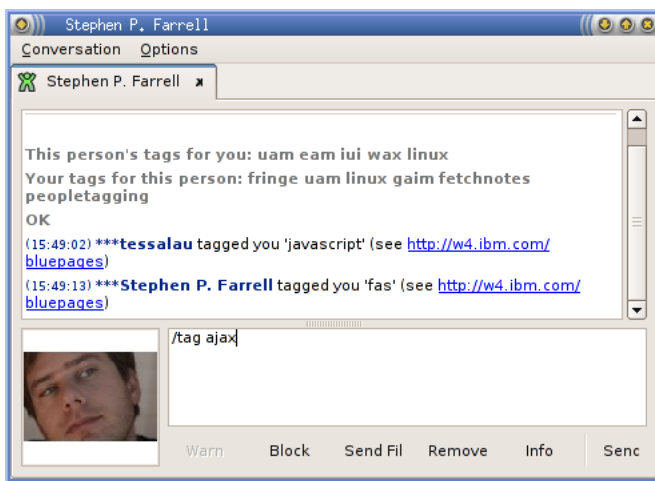


Figure 3: Tagging a chat buddy from within an instant messaging conversation

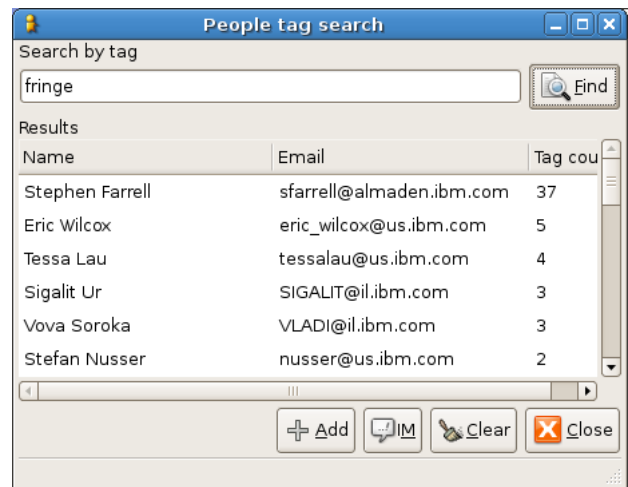


Figure 4: Searching for someone by tag in order to initiate a chat

One such integration point is instant messaging. We developed an extension for an instant messaging application that enables users to tag others, and view their tags, from within the context of a chat conversation (Figure 3). When a new conversation window is opened, the system displays the tags you have applied to your buddy, and the tags he/she has ap-

plied to you. During the conversation, the plugin provides commands to add a tag to the person, delete a tag, replace the set of tags with a new set, and display the person's tag cloud (list of all tags ascribed to this person by anyone, ordered by frequency). Changing someone's tags results in a message being generated in the chat conversation visible to both part-




Figure 2: Tag group view, showing the list of all the people company-wide who have been tagged with this tag. People are ordered by rank, with the more frequently-tagged people showing up first.




Figure 5: Tagging someone from within email

ners; this practice encourages users to spread awareness of people-tagging.

Our plugin also provides a tag-based buddy list: a dynamically-generated list that automatically contains the set of all people the user has tagged. Hovering over the name of each individual in the list brings up a tooltip with details about the contact, which we have enhanced to include the tags you have given to that person. We have also provided the ability to search for and initiate conversations with people based on tags. After typing a tag into a search box, the system displays a list of users matching that tag, and the user can click on a person to initiate a chat with them (Figure 4).

Another integration point is email. We have developed a plugin for the Thunderbird email client to integrate with the tagging system (Figure 5). Right-clicking on any email address displays the tags one has applied to this person, and provides a menu option to add new tags to this person.

A third integration point is web browsing. One of our colleagues has developed a plugin for the Firefox web browser called “Tommy!” that scans each web page for person identi-

fiers (typically an email address, a link to the employee directory, or a link to someone’s internal blog). Right-clicking on such a link causes a popup to appear with information about the person from Fringe and elsewhere. Tommy! includes a contact list that is synchronized with the tagging and friending data from Fringe. It supports searching for the user’s contacts by tag, or searching for people by tags given by anyone. Users can also tag other users from Tommy! and invite them to connect. This plugin enables a user to quickly discover information and assign tags to people they encounter while surfing the web.

Developer API

In addition to the web interface and the various plugins, we have provided a web service with a documented REST [3] API that enables developers to access and contribute tag data. All of the systems we’ve described depend on this service. Moreover, it has enabled the development of a handful of visualizations and mash-ups of tags with other systems. For example, some people have added a tagging widget to their blog pages to enable readers to tag them directly while reading their blog. Another created a social network visualization tool. A related project (a web-based mail client) is using the REST API to integrate tags into the email interface. In our opinion, this openness is essential to making people tags a ubiquitous feature in collaborative software, which in turn contributes to a large base of people tags.

USAGE STATISTICS

The people tagging functionality has been in use in our organization for approximately thirteen months, starting with an initial tagging-centric web application (described in a previous paper [2]) and more recently becoming integrated with our next-generation employee directory, Fringe.




Figure 6: Tags created per month

Currently the system holds 73,438 unique (tagger, tag, taggee) tuples. To illustrate system growth, Figure 6 shows the number of these tuples created each month. The spikes seem to coincide with new features (e.g., friending was added in December 2006), or announcements about Fringe to new communities. There are 10,658 unique normalized tags.

As of March 29, 2007, 3,462 distinct people have tagged at least one person, while 29,484 distinct people have been tagged by at least one person. On average, each person has tagged 12.10 other people (median 1, stdev 41.23, range 1-541). On the other hand, for any given taggee, on average only 1.42 people have tagged him or her (median 1, stdev 1.70, range 1-76). These statistics suggest that a small number of active users have contributed most of the tags. In fact, the top 10% (346 people) have tagged 85% of the taggees (25,026 people). This result supports our hypothesis that social tagging enables a small community to help maintain the directory as a whole.

Insofar as social networks are scale-free, we can expect this community effect to be robust both for much larger or much smaller numbers of users. The “long tail” of interests and expertise may also mean that infrequently used tags can be of high value to a small number of users. For example, a minority community of practice (e.g., designers among an organization of developers) may not tag as often as the majority community, but their tags may be more valuable for other members of that minority community. Future work will investigate this phenomenon in more depth.

AUGMENTING PROFILES WITH TAGS

We have investigated how tagging contributes to enhancing an employee’s profile. Each employee at our corporation can update their profile (in the online employee directory, which is widely used) at any time using a simple web-based application. However, not all employees have chosen to populate their profile with meaningful prose, nor does everyone keep their profile up to date with their most recent projects or expertise. Our hypothesis is that people-tagging, in which others contribute meaningful words to describe people they know in the organization, can help populate employee profiles with more timely information about their interests and expertise, and thus enable people to find other people more effectively within a large organization.

Methodology

We observed several differences between the tags used to characterize people and the prose in their profiles. First, taggers often conjoin multiple words with dashes to form a single tag, such as “it-specialist” or “websphere-portal”. These compound tags are unlikely to appear literally in plain text, but rather as phrases such as “IT specialist” or “Web-Sphere Portal”. Thus we opted to split all text (tags and prose) on non-alphanumeric characters, as well as lowercasing all strings, to facilitate comparisons across the two different systems.

Second, we noted that different word forms are often used to express the same information. For example, someone may be tagged as a “designer” or with the term “design”. As these different forms are conceptually similar, we used a stemmer [12] to canonicalize tags and words, and we base our comparison on the stemmed forms.

For clarity, we refer to the part of a tag that has been split on non-alphanumerics and stemmed as a *tag fragment*. After this preprocessing, our dataset contained 33889 unique (tagger, tag fragment, taggee) tuples.

Datasets

For each employee, we collected the words visible on their profile. The profile includes text that the employee provided, such as a biography or statement of interests. It also includes organizational information not written by the user, such as job title and department name. The profile also includes a limited amount of data drawn from internal data sources such as the titles of recent publications and patents. All of this text typically appears on each employee’s profile page in BluePages.

We denote the set of words on an employee’s profile as T_{BP} (assuming that each word is roughly equivalent to a tag). If $frag(t)$ denotes the set of fragments of word t , then let F_{BP} be the set of fragments generated from T_{BP} after tokenizing and stemming each word in T_{BP} . In other words, $F_{BP} = \bigcup_{t \in T_{BP}} frag(t)$.

To compare profiles with tags, we computed the set of tags with which each employee had been tagged (this included self-tags—tags the employee ascribed to him or herself). The resulting set we denote T_{in} , for incoming tags. Similarly, we also compute T_{out} , the set of outgoing tags this employee has used on himself or others. Let F_{in} and F_{out} be the corresponding sets of tag fragments generated from each of these tag sets.

Tags contribute information

Our analysis aimed to measure how much additional information is added to each person’s profile through the use of tags. For each person who had been tagged in Fringe, we calculated $F_{novel} = F_{in} \setminus F_{BP}$, that is, the set of tag fragments in their incoming set that were not present in their profile text. These fragments might represent new information or things a person is known for that they do not choose to present on their profile, but that might be useful to others in trying to locate expertise.

Of the 29,484 people that had incoming tag fragments, 24,552




Figure 7: Survey results on people tagging

(83%) of them had a non-empty F_{novel} , suggesting that the vast majority of people received tags that were not reflected on their profile. On average, including people whose tags did not contribute any novel words, people acquired 1.90 novel tag fragments via people-tagging (median 1, stddev 2.60, range 0-67). While it is difficult to interpret the meaning of tags without asking the tagger or having familiarity with the subject matter, a cursory scan of the novel tag fragments reveals internal project names, attributes such as `mentor` or `coach`, technical skills such as `eclips` or `ajax`, locations such as `frankfurt` and `nyc`, and roles such as `blogger`.

Though our data are still preliminary, we believe this evidence shows that people tags already provide distinctive information about a person beyond what that person has provided in their profile. And as people tagging becomes more ubiquitous, this value will continue to grow.

USER SURVEY

We crafted a web-based survey with 15 questions covering user satisfaction and proficiency with the Fringe employee directory, general demographics, and specific questions about awareness of and reactions to the people-tagging feature in Fringe. A survey request was distributed via email to thousands of people who had previously registered as “early adopters” in our corporation; these people had previously received announcements about the Fringe system being available for use, but have not necessarily used it. 129 people who received the invitation responded to the survey, and 63 of those respondents reported having used Fringe. Respondents came from all areas of our corporation, with the largest fraction (22%) representing our product divisions, and the next largest (15%) coming from our services division. In the analysis below, we consider only results from the 63 respondents who have used Fringe.

Figure 7 shows a summary of responses from the 63 respondents who had used Fringe. 68% claimed to be familiar with its tagging functionality. The majority of those familiar with tagging (74%) reported that they were either very satisfied or somewhat satisfied with the tagging feature; only 5% indicated that they were somewhat dissatisfied.

One survey question asked what respondents thought about having tags created by others show up on their profile. 64%

of the users liked it, 29% said it did not bother them, and 7% expressed their dislike for the idea.

In a follow-up question, the survey asked if people were comfortable with the tags they had received or if they wished to remove any of them. 86% indicated that they were comfortable with the tags they had received and only 14% expressed a desire to have tags removed. We would like to point out that this does not necessarily imply that these 14% of the people object to being tagged in general—some of the reasons that were mentioned in the comments section were “no longer relevant” or “remove tag I added”.

USER INTERVIEWS

In order to better understand how people used tags, we interviewed a smaller number of people in depth about their experiences with tagging people.

Interview methodology

We solicited 37 users who had the most incoming or outgoing tags in the system to participate in a 30-minute interview by phone. The first 19 who replied were chosen to participate in the interview. These people came from different parts of our business (only one from the research division), and several different countries.

The interview followed a semi-structured format. We developed a list of 10 questions to ask all participants, but we omitted some questions if they had already come up in the interview. The general format of the interview was to ask each user to examine their own profile and answer specific questions about what they saw. One team member asked questions while another took notes. These notes included answers to specific questions, transcriptions of volunteered information, and any compelling quotes. Quantitative results were drawn from these notes by tallying responses. The questions are summarized in Table 1.

Our interview covered four main topics. First, we wanted to understand how people used tags, particularly whether they used tags for contact management. Second, we examined how people felt about the tags they were given. Were they accurate? Did they object to any of them? Third, we compared people’s tags with their BluePages profile text, and asked whether there were any objectionable tags they might not want appearing on their profile. Fourth, we questioned whether tags could be used as votes to rank experts in a particular topic. We also asked participants about the usefulness of private tags.

Tags for contact management

Our hypothesis was that people would use tags to organize their contacts, bookmarking them for later recall, or browsing tags to find people related to a particular area. In response to the question about how people use tags (Q6), all 19 interview respondents reported using tags to organize their contacts. 17 reported clicking on a tag to recall people, and about half mentioned that they found it useful to see their own tags while browsing others’ profiles, reminding them what they knew of the other person. One respondent mentioned “unexpected value” in his tagging, because he found associations between people he tagged that he might not have made other-

- Q1) Do your incoming tags adequately describe you?
 Q2) Are there any tags you wish to have removed, and if so, why?
 Q3) Does the biggest/boldest tag describe you adequately?
 Q4) Of the tags in your incoming cloud that do not appear in your BluePages profile, how many accurately describe you?
 Q5) Why didn't you add these tags to your profile?
 Q6) Explain why you tagged people with some of your highest-ranked outgoing tags. Have you ever used tags to find people?
 Q7) Rank your three tag clouds in order of how they characterize you.
 Q8) Click on the highest-ranked tag in your incoming cloud, and examine the set of people who have that tag. Do you feel the first few people are the most associated with that tag/topic? Anyone out of order? Anyone missing? Anyone you'd add?
 Q9) If you had an option to assign private tags, would that encourage you to add additional tags, and why?
 Q10) Are there any other experiences with tagging you'd like to share with us?

Table 1: Summary of interview questions

wise. Another described “literally surfing the org chart” after switching teams and tagging the people he thought he would want to find again in the future.

For the purpose of organization, users typically chose tags that reflected either projects or skills. For example, `p-vista`, `dogear` and `phpzero` are names of projects that respondents used to label people. `php` and `design` are examples of tags that represent skill areas. In addition, some respondents reported tagging by geography (`italy`, `slovakia`), organizational unit (`cio`, `research`), brand (`websphere`, `tivoli`), context in which the person was found (`blogger`), or interest (`photography`). While some of these tags duplicated information already in the profile—e.g., `geography`—the respondents explained they used them to distinguish *their* contacts from a larger group.

We also found much evidence that people use tags to create groups of people for the purposes of defining a community. Several people mentioned that they would share these tag groups with others in the community, perhaps by sending a link to the tag group in email, or by referencing it from their community's web page.

Tags representing self

The second part of our interview investigated the use of tags to describe oneself (Q1, Q3). Every person interviewed said that the tags they were given by others characterized them accurately. Some mentioned that they had interests that were not completely covered by tags, which is due in part to the fact that our community of taggers tends to be disproportionately interested in Web2.0, social networking, blogging, and other modern technologies.

We also asked our interview respondents to rank the three tag clouds that appeared on their profile in the order in which the tag clouds best described them (Q7). The first tag cloud

contained the tags others had used to describe them, the second contained the tags they had used to describe others, and the final one included tags from a social-bookmarking application known as Dogear [8]. We found that 11 of the 19 respondents thought the first tag cloud best represented them, 1 favored the second, and 5 favored the third. The other two found the first and third clouds equally useful.

Three interview respondents volunteered that the first cloud was their “mirror”—how others saw them. One interesting result was a respondent who thought that this first tag cloud was understandable but not how he would like to present himself (the most prominent terms were his organizational unit and job category). A respondent who chose the second tag cloud was a technical sales representative who had tagged his extended network with geography and brand. He explained, “I have access to 299 people in 25 different countries.” 11 users volunteered that the Dogear cloud showed their interests. However, 4 respondents said their interests were varied and do not describe them well. One put it, “it's just something that's interesting rather than something that's core.”

Tags for profile augmentation

In another portion of the interview, we compared each interviewee's incoming tags with the text that appeared on their BluePages profile, using the same stemming and tokenizing approach described above.

We asked our subjects about the tags they received that were not reflected on their BluePages profile (Q4, Q5), and received a variety of answers. 16 described their profile as out-of-date and many said they “should” update it. Some explained that they no longer used BluePages in favor of Fringe. One respondent put it this way: “seeing all that stuff in the expertise summary reminds me that I need to update my record.” He then listed which terms he'd remove and stated that importing the tags others had assigned to him would “certainly suffice for updating my profile.” Several others also reported that they would have liked to have these words on their profile but had not yet taken the time to do so themselves. This result supports our hypothesis that social tagging will be a useful way to populate employee profiles with information contributed by others.

Some people talked about the differences between their profile, which they perceived to be largely static (like a curriculum vitae, a summary of the person's last ten years of experience), versus their tags (which reflected their current interests, and might also include extracurricular activities). For example, one person was tagged with the term `blogger` indicating that he contributes to IBM's internal blogging community. While this bit of information might be useful to someone wanting to make a connection with him, he was reluctant to add that term to his profile because he does not believe he is an expert in blogging. Similarly, a respondent tagged `leader` and `organized` stated he would not put these terms in his profile. This result suggests that information appears in tags assigned by others that would not appear in users' profiles, even if they were diligent in maintaining them. One person who used the tag `tap-first-adopter` felt that people were proud to receive that tag because of their

enthusiasm for being on the cutting edge of software.

Tags provide the opportunity for third-party-generated content to appear on one's professional profile. We were concerned that people might object to the tags they were given, thereby precluding the use of tags as a mechanism for augmenting profiles. *None* of the respondents reported having been given any inappropriate tags (Q2). However, we did find some tags that seemed more playful in nature. For example, we found the use of *sassy*, *funny*, *needs-a-shave*, *cleverclogs*, and *evilgenius*. The respondents who received such tags were not bothered by them. One user even encouraged the use of such tags, explaining that he thinks exposing his personality online is important, and that that is part of the idea behind social networking. He reported that he would not tag himself with playful tags, yet he used them to tag others. Another respondent tagged people *no-picture-in-bluepages* in order to shame them into adding a photo to their profile. He reported success with this experiment.

Tags for ranking

We believe that tags can be used as "votes" to determine which people are most associated with a given topic. We asked the interview respondents to click on a tag that was prominent in their incoming tag cloud and evaluate the list of people displayed (Q8). The resulting list (e.g., Figure 2) is sorted by tag rank, with the most frequently-tagged person shown at the top. With this question we hoped to get active members of those communities to judge how well tag-based ranking performed. 18 of the 19 respondents felt that the list they examined was representative of the people most associated with the tag. 2 suggested they thought the ranking order was wrong for some people, 8 could cite people they thought were missing, and 5 commented about people they didn't know.

We were interested in how the interview respondents perceived the ranking. One explained, "The first two rows are the people who are the most like-minded, the creative team... the broadest visibility on the IBM intranet." Another agreed: "The people who are here are very visible, very active". Others found the ranked listings they looked at less useful. One said the ranking for *cio* seemed arbitrary and that the first person listed was the *former* CIO. Another chuckled and said about the tag *userexperience*: "People on this list are not at all representative. I don't find the ranked list that helpful. A lot of these people I didn't tag and I cannot relate to them. If they don't have user experience in their title I cannot understand why they are here." Another perceived a similar listing differently: "People I'm not aware of is a good thing, because if I wanted to find someone with a skill I'd find people here I'm not aware of." While these results are compelling, the tagging data to support ranking are limited. We look forward to studying this effect in further detail.

Private tags

Our system does not support the creation of private tags. We asked (Q9) whether they would use such a feature. Four respondents thought they would find private tags useful. Some immediately jumped to the conclusion that private tags would be useful for making negative assertions, but only one thought

this would be productive. Others suggested that private tags might be useful for terms that are meaningful to them but not to others.

DISCUSSION AND DESIGN IMPLICATIONS

Our user interviews revealed that the tagging feature was being used as intended for contact management. However, several interesting themes emerged from the interviews that provide important implications for the future design of people-tagging systems.

For example, our results show that utilization of tags for managing contacts is still quite rare (only 1% of our company has created at least one tag). We expect that the primary issue here is accessibility: most employees "live" in email and instant messaging, where the tagging features are currently unavailable¹⁰. To tag a contact, a user has to navigate to his profile in Fringe. In contrast to social bookmarking where a browser plugin is sufficient, people-tagging seems to require plugins for email, instant messaging, web browsing, the corporate directory, and so on, to meet its potential. The productivity gains of tagging someone in their profile and then being able to reuse that action when composing an email message, e.g., could be considerable.

We found it encouraging that the vast majority were pleased with the tags others assigned them. However, tag obsolescence will become a problem over time. Several people reported receiving tags for projects with which they are no longer associated, and were concerned that these tags would appear on their profile indefinitely. In response, we are considering mechanisms to enable tags to decay over time.

We observed that most people tended to be extremely aware of tagging as a social activity. People think about how others will react to the tags they give. One person reported that he refrained from tagging one of his colleagues with the name of a project he had previously worked on. His rationale was that the recipient might not appreciate being contacted about the defunct project, and therefore he did not feel it was appropriate to publicize that information, even though it would have been personally useful to him to remember that fact about his colleague. In order to amplify the social considerations, we are considering notifying users when they are tagged, perhaps on an opt-in basis by the tagger.

The interview results suggested that incidents of people wanting to remove tags will be rare. In addition, as administrators of this system, we have yet to receive a single request to remove a tag from someone's profile. Nevertheless, incidents will undoubtedly occur, especially as usage grows beyond the self-selected early adopter crowd. It has yet to be seen whether a social back-channel is really sufficient to address this issue, or if some other mechanism will be required.

Another issue (which we expected to be more widespread than we found) is tag naming collisions. One user was working on a project codenamed "Vista", which is also the name of an operating system. Because he was concerned that people would erroneously believe he was associated with the operating system, he chose to use the tag *p-vista*, and asked

¹⁰The plugins we wrote for this purpose are undeployed prototypes.

his fellow project members to do the same. While interface features such as tag auto-completion provide some support for choosing appropriate tag names, this problem is still far from solved. It is likely that this problem will grow worse as the system is used more. A tagging system might need to support social scoping of tags so that small teams can use the same tag without interfering with each other.

CONCLUSION AND FUTURE WORK

We have described the use of people tagging in our Fringe enhanced employee directory. We have shown that people use tags for personal benefit, to label and recall other people, much in the same way they socially bookmark web pages. Our user studies have confirmed that people do perceive valuable benefits from using tags, such as being able to see all the people associated with a project, or locate experts in a particular area.

One of the potential concerns about tagging people (as opposed to bookmarks or other resources) is that people might be offended by a tag they have received. However, none of the people we surveyed reported any cases of inappropriate or offensive tag use; in fact, several people reported that they expected to use social back-channels (discussing the tag in question with the tagger) to solve the problem. We are optimistic that people tagging will become an important tool for people to manage their relationships with other people.

In addition to the personal benefits of tagging, we have found that tagging provides an intriguing side benefit: it enables an active user community to help maintain the employee directory as a whole. We believe tagging distributes the work of maintaining one's profile out to the people who stand to benefit most from having one's profile current. This raises the question of whether users should be given more power to maintain others' profiles.

Another possibility we are intrigued by is the use of tags as a ranking function for finding people. Our interview results indicate that our current system of using tags as "votes" is moderately successful. Future research will investigate better ranking functions that combine tag-based voting with other metrics of expertise. Furthermore, a common concern about "expertise location" systems is that, were they to work well, the experts would be overwhelmed with questions they had no interest in answering. It seems that the people who rise to the top in the tag-based ranking are the people who talk about the topic most in blogs and other forums. Perhaps tag-based ranking brings out the "hubs" rather than the subject matter experts, and perhaps is more useful for finding the right person. We plan to investigate this phenomenon in more depth as the usage of tagging in our corporation increases.

We are particularly interested in the use of tags to create communities, and enhancing our tools to better support that use case. For example, an existing community membership tool is considered onerous but is still necessary for managing access control. Can people-tags be used to manage access control?

Another direction of future work involves comparing the tagging of people in Fringe with tagging behaviors in other en-

terprise-oriented tagging systems, such as webpages and resources [8] and activities and their components [10].

ACKNOWLEDGMENTS

We would like to thank Helder Luz for his excellent Tommy! Firefox plugin and his help in debugging the Fringe Contacts API. We thank Effie Seiberg and Elizabeth Palmer for composing and conducting the first Fringe survey.

REFERENCES

1. T. Erickson and W.A. Kellogg. Social translucence: An approach to designing systems that mesh with social processes. *Transactions on Computer-Human Interaction*, 7(1):59–83, 2000.
2. Stephen Farrell and Tessa Lau. Fringe contacts: People-tagging for the enterprise. In *Workshop on Collaborative Web Tagging, WWW 2006*, 2006.
3. R.T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, Irvine, 2000.
4. E. Goffman. *Presentation of Self in Everyday Life*. Doubleday, New York, 1956.
5. Scott Golder and Bernardo A. Huberman. Usage patterns of collaborative tagging systems. *Journal of Information Science*, 32(2):198–208, 2006.
6. A. John and D. Seligmann. Collaborative tagging and expertise in the enterprise. In *Workshop on Collaborative Web Tagging, WWW 2006*, 2006.
7. C. Marlow, M. Naaman, d. boyd, and M. Davis. Position Paper, Tagging, Taxonomy, Article, ToRead. In *Workshop on Collaborative Web Tagging, WWW 2006*, 2006.
8. D.R. Millen, J. Feinberg, and B. Kerr. Dogear: Social bookmarking in the enterprise. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 111–120, New York, NY, USA, 2006. ACM Press.
9. Audris Mockus and James D. Herbsleb. Expertise browser: a quantitative approach to identifying expertise. In *ICSE '02: Proceedings of the 24th International Conference on Software Engineering*, pages 503–512, 2002.
10. M. Moore, M. Estrada, T. Finley, M.J. Muller, and W. Geyer. Next generation activity-centric computing. *Demonstration at CSCW 2006*, November 2006.
11. B.A. Nardi, S. Whittaker, E. Isaacs, M. Creech, J. Johnson, and J. Hainsworth. Integrating communication and information through ContactMap. *Commun. ACM*, 45(4):89–95, 2002.
12. M.F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
13. Z. Xu, Y. Fu, J. Mao, and D. Su. Towards the semantic web: Collaborative tag suggestions. In *Workshop on Collaborative Web Tagging, WWW 2006*, 2006.