RJ10384 (A0606-027) June 30, 2006
Computer Science

IBM Research Report

Fringe Contacts: People-Tagging for the Enterprise

Stephen Farrell, Tessa Lau
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099

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

Fringe Contacts: People-Tagging for the Enterprise

Stephen Farrell and Tessa Lau
IBM Almaden Research Center
650 Harry Road
San Jose, California
+1 408-927-{1856, 1901}
{sfarrell@almaden, tessalau@us}.ibm.com

ABSTRACT
Tagging has arisen as way to enable users to contribute to a loose taxonomy characterizing web pages, pictures, products and other things. We propose tagging *people* in order to help individuals keep track of each other while contributing to a loose characterization of their friends and colleagues. "Fringe Contacts" is a reference system designed to test whether people-tagging is a viable and useful approach. It includes both user and programmatic interfaces to tagging functions. The latter is included to enable integration with other collaborative applications such as email or instant messaging. Some client plugins have been developed, and preliminary usage data are encouraging. We think that some characteristics of the enterprise —a clear notion of identity, "professional" environment, and existing need to classify people by their skills and projects—may be necessary ingredients for people-tagging to work.

Categories and Subject Descriptors
H.5.3 [Group and Organization Interfaces]: Collaborative Computing

General Terms
Design, Experimentation, Human Factors

Keywords
Tagging, Social Networks, Relationship-Oriented Computing

1.INTRODUCTION
Tagging has recently acquired popularity as a lightweight and flexible approach to classifying information. Tagging enables individuals to use whatever terms they think are appropriate to describe or help them recall a resource without the burden of selecting a category from a known taxonomy. It has been applied in a variety of applications ranging from desktop applications for organizing photos (F-Spot [5]) to web email systems (Gmail [6]). Tagging becomes most compelling when it is used in a collaborative environment, and tags from different people can be aggregated and combined. This approach has been used to manage bookmarks (Del.icio.us), images (Flickr) and products (Amazon.com).

We want to know if this same approach can be used to characterize *people*, particularly in a large enterprise where maintaining and forming relationships are routine parts of the lives of thousands of people. This paper introduces the ideas and motivations behind people-tagging, describes our research implementation called "Fringe Contacts", and provides a preliminary discussion of the technical and social issues raised.

2.TAGGING PEOPLE
Like many others, we have observed the trend toward applications that enable users, motivated by their own interests, to provide information that benefits a larger community. Del.icio.us, one such application, enables users to keep track of and share web pages while at the same time contributing to an emergent taxonomy (or "folksonomy" [1], [2]) that provides benefits even to users who never bookmark anything themselves. Our ongoing Fringe project has been exploring tools for people to find, learn about, and keep track of each other in enterprise environments in order to improve the effectiveness and reduce the cost of forming and maintaining professional relationships. We see some compelling parallels between social bookmarking and relationship management. Specifically, as with bookmarks, people presently keep track of colleagues using address books, buddy lists and other personal information management tools. Also, while the organization chart and employee directory provide key data about individuals, there is a persistent need to show more about people and how they are connected to others. Current efforts in enterprise directory development emphasize developing a fixed classification scheme [3], similar to the approach taken by web taxonomies such as Yahoo! Directory and DMoz. Our idea is that the same tagging mechanism and emergent "folksonomy" that works for bookmarking web pages can be adapted to work for bookmarking people in the enterprise, and that it will have the side-effect of providing rough descriptions of people's skills, roles, and projects in the form of a "tag cloud". Similarly, the frequency of a tag should say something of its strength, and so someone tagged with "java" 78 times is likely to know more (or, more accurately, to have a greater *reputation* for knowing more) than someone tagged 4 times with the same tag.

Tagging people is also a way to capture limited social-network data. One can see a little bit about who someone knows and who knows them and, more usefully, who the user and another individual have in common. Putting the descriptive and social network aspects together could yield even more compelling results. Seeing, for example, that someone has been tagged "ruby" by someone you know says that your colleague knows this person, that your colleague knows enough about "ruby" to make this judgment, that this person knows about "ruby", and that someone thinks they are knowledgeable about it. This tag enhances the reputation of the person who received this tag, particularly if they have received it from others as well. Moreover, the social context might reduce the ambiguity that would otherwise exist about whether this was "ruby" the programming language or "ruby" the mineral, etc.

While there are precedents for people-tagging including Tagalag [4], we believe that the enterprise environment has some characteristics that make people-tagging more viable than on the Internet. For instance, the set of everyone in the organization is typically known, and each person has a unique identifier such as

an email address; the tagging system can thus be pre-populated with everyone in the organization, without people having to explicitly add new members. The existing employee directory can be mined for initial data such as real name, location, phone number, job title, and in many cases a picture. In our organization, intranet access is authenticated, so that intranet web sites such as Fringe Contacts can verify your identity without requiring you to create an account on the system and log in each time. This lack of anonymity and the professional context, in which people are accountable to their management, reduce the likelihood that users will abuse the system. (However, as discussed below, they do not preclude social issues from arising).

Our goal of optimizing relationships also includes personal relationship-management enhancements, and we think that Fringe Contacts can provide some useful services toward that end. For instance, as a contact-list back end, it can unify buddy lists and address books across applications. By tagging someone who you looked up in the enterprise directory, you can then send an instant message or email to the person by simply typing that tag (or just a part of their name) from your email or instant messaging client. Likewise, if you tag someone you received mail from, then look up their co-worker a week later, the enterprise directory can highlight this common connection. People can also use tags as lightweight mailing lists or groups. To support all of these functions, we envision a centralized address book service that interacts with plug-ins to existing productivity and collaboration applications.

Designing and developing Fringe Contacts had led us to consider some hypotheses testable by further research:

1. Tagging is an effective way to organize contacts
2. "Social" tags can inform others about someone's interests, skills and expertise
3. People-tagging benefits from properties unique to the enterprise including a pre-populated directory
4. A tagging service can reduce the cost of relationship building and maintenance with an integrated experience across applications

3. IMPLEMENTATION

In order to learn more about people-tagging, we created a web application with an HTML-based user interface and programming interfaces implemented as web services using the REST [7] design pattern. The user interface is not expected to be the primary way that people interact with tags, but is helpful as a reference and starting point. The web services enable developers to bring people-tagging into the context of other applications that people routinely use such as their web browser, calendar, email and messaging applications. Of course, data is shared between these systems and, in a sense, the web user interface is just another client of the underlying tag system.

3.1 Representation

We model tags as properties of the directed edge that represents one half of a relationship between two people. So if Randolf tags James "manager", that suggests that James is "manager" to Randolf, but not vice-versa. James might choose to reciprocate this tag with "direct-report". Internally, the directed edges are represented as an ordered pair of people (the *user* and *target*), an id, and created, modified and accessed timestamps. Each edge can have one or more tags associated with it (setting the number of tags to zero removes the edge, a decision we may revisit). The direction of the edge also is linked to the authentication mechanism: users may only create and update tags on edges in which they are the *user*. It is possible to pivot on the *user*, *target*, or any of the tags. So, for example, one can efficiently retrieve everyone tagged by Randolf, everyone who has tagged Randolf with "colleague", every tag Randolf has used, and so on. The user and programming interfaces into the tagging system work by issuing these queries and formatting the results.

3.2 Web User Interface

The web-based UI enables users to look up other employees, associate tags with them, and see what tags they have used and have been used to describe them. Any user can create tags and associate them with any other user. For example, one of the authors has tagged 143 people with 100 tags ranging from "fringe", the name of this project, to "jim", the nickname of a former boss. All tags are public, so anyone can see who he has tagged with each of these tags, and, given any person, what tags have been associated with him or her. Next to the profile of the person is a list of related people. This list is computed by finding people with the most tags in common with the person currently being viewed. In contrast to tagging bookmarks and products, tags in Fringe Contacts can be reciprocated: the people a user tags may tag that user back. As a result, it is necessary to distinguish between "incoming" (left column) and "outgoing" (right column) tags for each person (see Figure 1).

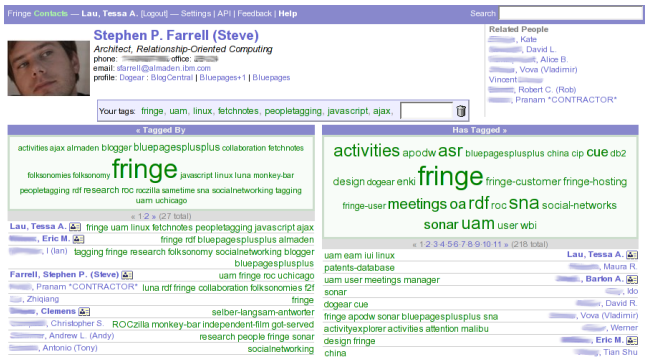

Figure 1. Fringe Contacts Web User Interface

The tags used on and by a person are aggregated into "tag clouds". These collections of tags are sorted alphabetically, but typographically modulated based on the frequency of a tag: more frequently-used tags are displayed in a larger font. The clouds also serve as filters: clicking on one filters the list below to show only those people who have either used the tag on this person, or that the user has described with the tag. One can follow a second link to get a tag-centric view from which they can see everyone who has used the tag and everyone on whom the tag was used. In this view, the most frequently co-occurring tags are represented as another tag cloud.

The web interface has two interactive features that make it easier to create tags and find people who have been assigned tags. The "tagger" widget, inspired by Amazon's tagging function, enables users to create multiple tags without a page refresh, and visually distinguishes them as formatted, comma-delimited text, not a web text input element (see Figure 2). The "search" widget is used to find people. If a single term is entered, it checks if that term is part of any name the user has already tagged, or is itself a tag. In either case, the results are brought up as a sorted list. If a name was matched, then the results are sorted by recency. If a tag is matched, then the people the user tagged with that tag

appear first, sorted by recency, followed by those others have tagged with that tag, sorted by the number of times that tag was used on them. If both a tag and name part are matched, the name matches appear first (see Figure 3).

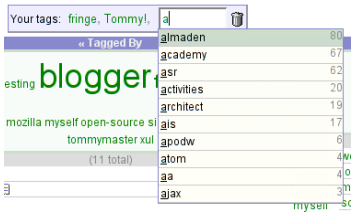

Figure 2. Tagger widget, including auto-complete on tags

Figure 3. Personalized, tag-based auto-completion

3.3 REST API

We believe that pervasiveness of tags is a prerequisite to Fringe Contacts being truly useful. Thus, in addition to the user interface described, Fringe Contacts also provides APIs that enable developers to provide a person-tagging capability in other applications. For example, if a user receives an instant message from an unknown colleague, and they chat for a few minutes about the Ruby programming language, then the user should be able to tag her colleague "ruby" so she will be able to find him in the future. We believe it is important that she be able to tag her colleague *in context*, ideally by tagging him directly from the chat application. Likewise, we hope to enable person tagging in any other context in which users encounter other people. Tags can also be used to identify *ad hoc* groups; for example, a user could send an instant message to everyone she has tagged with the term "lunch-gang", or filter her email inbox to only show messages from people tagged as "fringe".

The API follows the REST style, and uses simply constructed GET URLs to fetch data, and POST requests to update or modify data. For example, to find all the the people one of the authors has tagged with "fringe":

```
http://fringecontacts.example.com/api/
                    ?user=sfarrell@almaden.ibm.com&tags=fringe
```

likewise to find everyone who has tagged the author with "fringe":

```
http://fringecontacts.example.com/api/
                    ?target=sfarrell@almaden.ibm.com&tags=fringe
```

The result of these requests is an XML document such as:

```
<contacts target="sfarrell@almaden.ibm.com" tags="fringe">
          <contact>
                    <name>Jose Saramago</name>
                    <email>jose@example.ibm.com</email>
                    <tags>research people fringe</tags>
          </contact>
          ...
</contacts>
```

The available APIs include:

- **tags** -- shows tags one user has associated with another
- **tagged** -- all people tagged with a tag
- **contacts** -- all contacts of a person
- **cloud** -- all tags used by or on a person
- **tag** -- append, delete or set tags
- **stats** -- how frequently a tag has been used

Each API that returns results allows the developer to specify an offset and a limit, enabling navigation of large amounts of data.

3.4 Plugins & Integration

Several extensions and plugins have already been developed using the Fringe Contacts API. These plugins include a bookmarklet to enable integration with other web pages, a next-generation employee directory, a Firefox plugin that discovers related information to person links, and a Gaim plugin to integrate with instant messaging. As critical plugins for other instant messaging and email applications have not been developed, we consider this experiment to be in the early stages.

3.4.1 *Bookmarklet*

First, we provide a simple bookmarklet[1] that detects all references to people in the current web page identified either by email address or by a link that includes their email address. Each reference to a person the user has tagged is highlighted with a yellow background, and hovering over the reference pops up a tooltip showing the tags the user has applied to this person. This bookmarklet enables users to quickly "mark up" a web page to highlight all references to people in their social network. It works by loading a secondary script parameterized using GET request by all of the person-identifiers found in the document. This secondary script is generated based on these parameters, and performs the calls to annotate the original web page. This bookmarklet is a proof-of-concept, and we plan an enhanced version of this bookmarklet that enables users to update their tags.

3.4.2 *Fringe3 Directory*

Fringe3 is our latest revision of an advanced directory system. The tagging system integration is intended to help it to achieve the goals of showing more about people inside the organization and making it easier to keep track of colleagues (see Figure 4). Tagging is compelling in this environment, since it enables all users to contribute to the characterization and organization of any employee.

Fringe3 communicates with the tagging service through a number of the APIs listed above. It queries for all of the tags one user has applied to another**,** and retrieves the tag "cloud" for each user. Moreover, it provides a view of everyone that the user has tagged, as well as views of everyone tagged by a certain tag. In addition, the auto-complete widget is exported into this interface, enabling rapid lookup of tagged people. One interesting function is that when an authenticated user accesses the directory, the names of everyone they have tagged are visually highlighted on directory pages. Thus, as they are browsing through the directory, they can quickly see who has been previously tagged. Lacking from the current implementation is an indication of who assigned each tag (hovering over a tag in the "cloud" shows only the frequency

[1] A web browser bookmark that invokes Javascript code

with which the tag was used), and a way to see who someone's contacts are.

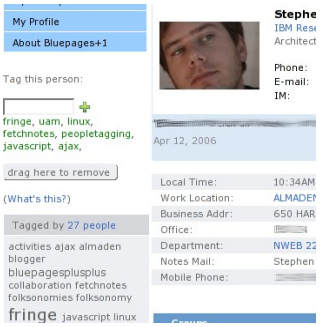

Figure 4. Part of Fringe3 corporate directory, a client of Fringe Contacts API

3.4.3 Tommy!

Tommy! is a Firefox extension created by a colleague of ours for internal IBM use. It works by augmenting the right-click menu for web links by checking for person-identifiers and, when present, providing a wide variety of data about that person gathered from internal databases. The author has extended Tommy! to support Fringe Contacts and, indeed, helped develop our API. It enables users to see and update their tags, as well as see who has tagged someone, and who that person has tagged (see Figure 5).

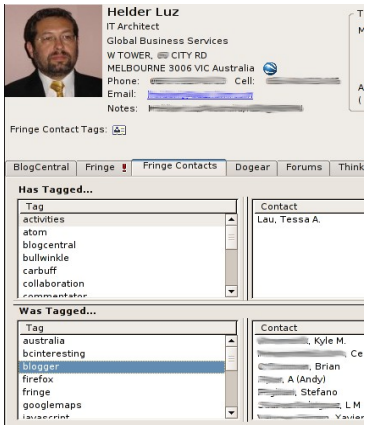

Figure 5. Tommy!, a Firefox extension and client of Fringe Contacts API

3.4.4 Integration with instant messaging

We have developed a plugin for the Gaim instant messaging application that enables users to tag others, and view their tags, from within the context of a chat conversation (Figure 6). When a new conversation window is opened, the system displays the tags you have applied to your buddy, and the tags he/she has applied 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 applied 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 partners; this practice encourages Fringe Contacts users to spread awareness of people-tagging.

The plugin also provides a tag-based buddylist. The system retrieves the set of people the user has tagged, and displays them as a dynamically-generated buddy list. Hovering over the name of each individual in the list brings up a tooltip with details about the contact, including your tags for that person. We also provide the ability to search for and initiate conversations with people based on tags.

We hope that providing the ability to tag people in context lowers the barrier to people-tagging, and results in increased usage of the system. Efforts are currently underway to port this functionality to other instant messaging clients in use within our organization.

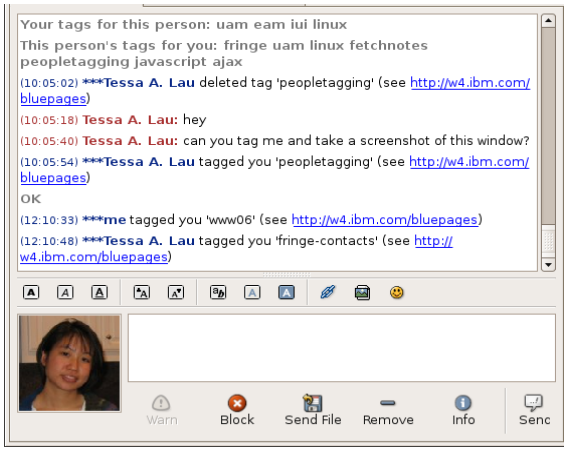

Figure 6. Fringe Contacts integrated into Gaim instant messenger client

4. EARLY RESULTS & DISCUSSION

While we are not yet in a position to test our hypotheses about people tagging in the enterprise, we do have some initial reactions.

4.1.1 Tags help individuals organize contacts

Our first hypotheses is that tagging is an effective way to keep track of and manage contacts. Although Fringe Contacts has only been deployed within our organization for approximately four months and no plug-ins have been written for the most widely-used collaborative applications – email and instant messaging – we have observed that some early adopters have started to use it. 304 people have created tags, entering a total of 1387 tags comprising 2937 person-person edges. The most popular tags were those generated by important "buddy lists" from instant messenger clients. The most popular tags (and frequency) include "blogger" (312), "sametime" (123) and "research" (89). Many of those tags were automatically applied when people imported their IM buddylists into Fringe Contacts. Other popular tags include "academy" (68), "cio" (63), "collaboration" (53) and "km" (46). Collaboration and km (knowledge management) refer to the interests of people using the tagging system, while academy and cio refer to membership in various IBM organizations. Tags such as "architect" (42), "design" (33), "de" (26, an IBM Distinguished Engineer) and "inventor" (15) speak more about the skills of individuals. Clearly more data will be required for a proper analysis.

Future research will go into understanding how effectively people can use tags as a way to keep track of their colleagues and how much efficiency they gain from that endeavor. We cannot say much about these yet, but we have received some favorable comments from supporters and early-adopters including, "I think this level of interactivity is going to get people to tag each other with wild abandon and should revolutionize our view of social

software within IBM", and "...this is an awesome glimpse of the future."

4.1.2 *Tags inform about social network and expertise*

Our second hypothesis is that tagging can provide information about individuals and their relationships to others inside the organization. In particular, we think that the tag cloud can provide information about someone's skills and interests, and that the people most associated with tags are somehow "important" relative to that tag. For example, the tags assigned to Helder (see Figure 5) indicate "blogger" and "tommy" are the most important terms characterizing him, and the others fill out some details including his location ("australia"), and some technologies related to Tommy! including "mozilla", "firefox", "javascript" and "xul". Helder is well known in the internal blogging community for his blog and for developing Tommy!, so these tags are appropriate. It should be noted, however, that they reflect his reputation among the early adopters of Fringe Contacts, and in fact say little about his primary job responsibilities. It is not surprising that looking at the tag "tommy" would show that Helder is the most important person relative to that topic, or that he is one of the most important people relative to the tag "blogger". What is surprising is that this information emerged as a result of voluntary action taken by people across the organization without any coordination.

There are also some potential problems with people-tagging that appear to stem from their very nature. Some users have expressed misgivings about others being able to type something that appears on their profile. In fact, we have had one incident so far that was relayed to the authors. This incident involved three people: Alice, Bob and Chuck (not original names). Alice knows Bob and Chuck well, but Bob and Chuck are newly acquainted. Bob used a few tags including "self-involved" to characterize Chuck. Chuck, not surprisingly, was a little put off and asked Alice if she knew why Bob had done this. Alice said she did not, but contacted Bob on the side. Bob's response made it clear that he was joking and, upon finding that it bothered Chuck, immediately removed the tag. This example is interesting because it illustrates, on the one hand, a big concern of tagging people—that they may not be pleased with the tags that others assign to them—and, on the other hand, how the openness (and lack of anonymity) of the system enabled social forces to be employed to correct the situation. It is suggestive, however, of bigger problems that might arise. For example, what if Bob had been unwilling to remove this tag? Or, what if Chuck had been too timid to approach Alice, or Alice unwilling to approach Bob? How these problems do or do not resolve themselves will be interesting to observe in the future.

Users have also been concerned about the permanence of tags. For example, what if many people have tagged them with the name of one project or team, and they move on to another. Should there be a mechanism by which tags decay over time? Will social forces be enough to correct these inaccuracies, or should a user be able to remove tags that others have assigned to them?

4.1.3 Tagging Benefits from Enterprise Environment

The enterprise environment provides comprehensive employee records, places constraints on who can use the tagging system, and the professional environment and lack of anonymity restrains behavior. In fact, the resolutions to both of the social problems referred to in the previous section depended on authenticated users who know each other and operate within a community. We have not tested whether these constraints are necessary or coincidental. Some possible areas of future research might be to try to enable person-tagging on the Internet in a way that leverages reputation to weight tags differently. Alternatively, we could see how enabling people to post tags anonymously would affect the social dynamics inside the enterprise.

4.1.4 Pervasive Tagging

The motivation for the tagging-as-service implementation is to enable people-tagging to be integrated into many existing applications with minimal effort. We think the few plugins developed already suggest that this approach is viable. It is not until users can tag others from any electronic context in which they encounter people (email, instant messaging, web browsing, teamrooms, etc.) that we will see how this impacts user behavior.

Another direction for future work is to compare people-tagging against existing collaborative tagging practice, such as social bookmarking or photo tagging. Do people use the same tags to describe other people as they do web pages or photographs? Does the fact that tagging can be reciprocated (I tag you, you tag me) lead to different tagging behavior?

In conclusion, we believe that people-tagging in an enterprise can provide many benefits, such as helping people maintain professional relationships, improving expertise location, and explicitly capturing the social network already present in an organization. However, our experience has shown that there are potential pitfalls in introducing this functionality into an enterprise. Thus our research goal is to investigate the challenges, both technological and social, of bringing people-tagging to the enterprise, and helping people maintain their professional relationships more effectively.

5. ACKNOWLEDGEMENTS

Helder Luz created the Tommy! plugin and helped debug the Fringe Contacts API. Eric Wilcox designed the interface for Fringe3 which integrated a tagging widget and inspired this project. Jon Feinberg's work on Dogear set the standard for lightweight, REST-based tagging services deployed within our organization (IBM).

6. REFERENCES

[1] Shirky, C. "Folksonomy" (August 2004), http://www.corante.com/many/archives/2004/08/25/folksonomy.php

[2] Vanderwal, T. (2005). "Off the top: Folksonomy Entries", http://www.vanderwal.net/random/category.php?cat=153 Visited January 28, 2006

[3] Voelker, M. "Optimizing the Human Supply Chain" (January 2006), http://www.intelligententerprise.com/print_article.jhtml?articleID=175002433

[4] "Tagalag: Manage your identity". http://www.tagalag.com/

[5] F-Spot. http://www.gnome.org/projects/f-spot/

[6] GMail. http://www.gmail.com

[7] Fielding, R. (2000) "REST - Representational State Transfer". http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm