Exploring the Adoption, Utility, and Social Influences of Social Bookmarking in a Corporate Environment

Laurie E. Damianos, Donna Cuomo, John Griffith, David M. Hirst, James Smallwood
The MITRE Corporation
{laurie, dcuomo, griffith, dhirst, smallwood}@mitre.org

Abstract

This paper describes an ongoing technology investigation to assess the value and utility of social bookmarking on a corporate intranet. We hypothesize that social bookmarking would be useful in this environment for resource management, information sharing and discovery, expert finding, and social networking. We discuss features of the prototype system deployed and early analysis of findings on adoption, usage, and social influences. We conclude with our challenges and plans for future development and integration into the enterprise.

1. Introduction

Social bookmarking is one of the latest popular social software technologies quickly gaining critical mass on the Internet. Other so-called "pop-tech" historically has included email, social networking sites, blogs, and wikis. *del.icio.us* [1] is one of the more popular, free-use tools for tagging and sharing web resources.

Traditionally, people have stored favorites or bookmarks in a browser client, such as Internet Explorer or Firefox. Browser bookmarks are displayed as a list of hyperlinked titles that can be edited, re-ordered, or filed within a hierarchical folder structure. These bookmarks are accessible only through the original browser and computer used to store them, and there is no direct way to share bookmarked resources with other people.

Social bookmarking differs from traditional bookmarking in several very critical ways. First, bookmarks can be annotated with user-specified metadata, including meaningful keywords which facilitate retrieval and act as memory jogs to bookmarked content. People can also add a textual description or comments to the bookmark. There is no hierarchical structure to social bookmarking; organization emerges from non-mutually exclusive categorization via tagging. Users can retrieve bookmarks by one or more tags (or title or description) without having to search nested folders. And, since bookmarks are stored in a central repository, bookmark collections are viewable in any browser on any Internet-accessible device. The "social" aspect of social bookmarking allows people to share their resources with others, explore the tag space, and discover virtual communities of others interested in the same topics.

The MITRE Corporation is a not-for-profit organization with expertise in systems engineering, information technology, operational concepts, and enterprise modernization. In addition to managing Federally Funded Research and Development Centers, MITRE supports its own independent technology research and application development for solving sponsors' near-term and future problems. MITRE has approximately six thousand scientists, engineers and support specialists who work on hundreds of different projects across various sponsors and numerous domains. Because of the high level of technical, operational, and domain knowledge required, staff often seek out and consult with MITRE experts on particular topics or seek assistance from in-house librarians in gathering resources. Much of MITRE's work involves collaboration across time and space, requiring virtual teams to share these resources and relevant research internally through the use of Listservs, technical exchange meetings, wikis, websites, and collaborative spaces.

The concept of social software maps well to our overall collaboration and cross-corporate information sharing goals, and with many MITRE employees already using popular Internet tagging tools such as *del.icio.us* and *Connotea* [2] for both business and personal use, we saw potential benefits to bringing in a social bookmarking tool to our closed, corporate environment. We believed social bookmarking would complement our current techniques for sharing corporate knowledge and help to expand social networks. Although there are currently no known

In Proceedings of the Hawaii International Conference on Systems Sciences (HICSS-40), January, 2007.
Approved for Public Release; Distribution Unlimited.
MITRE Case # 06-1199

commercial products available, BusinessWeek reports, "companies are figuring out ways to take advantage of [the social bookmarking] phenomenon" [3]. In fact, IBM is also experimenting with its own version of an internal social bookmarking tool, *dogear* [4, 5].

Using the MITRE Corporation as a case study, we describe an ongoing technology exploration to assess the value and utility of social bookmarking on a corporate intranet. We hypothesized that such technology would be useful to a corporate environment for resource management, information sharing and discovery, expert finding, and social networking.

Specifically, we would like to leverage social bookmarking to:

- provide teams with a place to share resources
- help form and support social networks around interest areas
- feed expertise finding & user profiling

The remainder of the paper describes our internal social bookmarking pilot, marketing strategies, observations on initial use, some issues and challenges, and plans for the future.

2. Pilot study and prototype description

We identified and interviewed several social tagging-savvy employees to get a sense of what they liked about their favorite tagging tool (e.g., *del.icio.us*, *Connotea*, *etc.*). After looking at the functionality and flexibility of several open source tools [6, 7], we decided to base our prototype on *Scuttle* [8], which shared many features of popular *del.icio.us*, was simple to install, and was easily modifiable. We adapted the software for use in our corporate environment (e.g., changed the look and feel), extended its capabilities by adding new features and integrating it with other corporate applications and resources (e.g., email, corporate directory), and promoted the sense of community by incorporating photographs of users and information on their organizational affiliation. In keeping with the popular convention of naming such tagging tools non-sensical or whimsical names, we dubbed our prototype *onomi* (rhymes with folksonomy, the term coined by Thomas Vander Wal [9]).

Figure 1 *onomi* home page

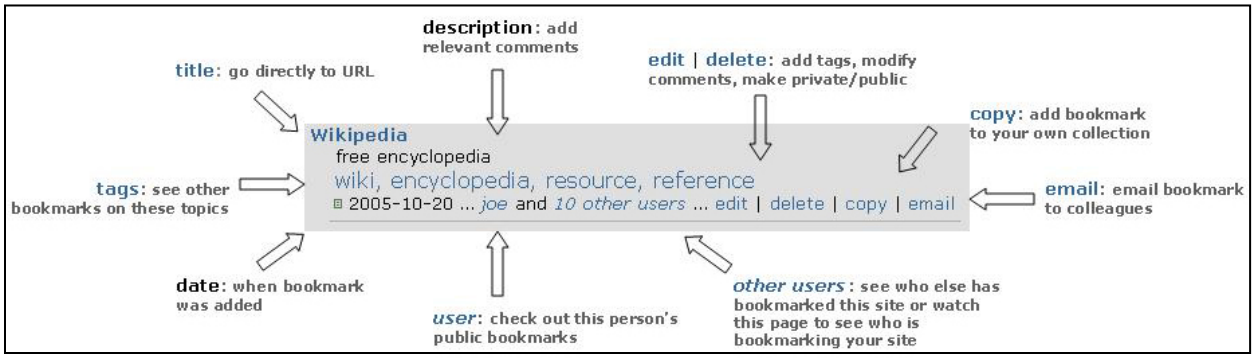

Figure 2 Anatomy of a bookmark

Figure 3 View of bookmarks shared by two users, with pop-up user information and a related user

We introduced *onomi* step-wise into the corporate community, first with selected sets of early adopters who were willing to test the prototype, populate it with useful resources, and provide us with iterative feedback throughout the development cycle. Some of these early adopters included our librarians who were already in the practice of gathering and distributing information but had no standard method of storing or managing pointers to these resources which they were frequently called upon to retrieve again for other MITRE customers. Through word-of-mouth and a series of low-key announcements and informal demonstrations, *onomi*'s user base began to grow. Our original plan was to pilot the social bookmarking system for six months, assess its utility, and then re-implement if successful.

The *onomi* home page is depicted in Figure 1. The page was designed to attract new users by highlighting the dynamics of the system: which resources people were adding (A), recent hot topics (B), popular topics over time (C), what people were looking for (D), and who was using the system (E). Recently-posted public bookmarks were displayed in reverse order, with the most recent posting at the top (Fig. 1, A). Users could set up an RSS feed of recent bookmarks (or specific collections and topics) to keep abreast of the company's expanding resources. (Bookmarks were also displayed by popularity on the *pop-onomi* page.) Each bookmark was displayed with associated metadata (see example in Figure 2), consisting of an editable title (hyperlinked to the bookmarked resource), an optional description, optional user-defined tags, post date, and owner name. The number of *other users* indicated how many additional users had bookmarked the same URL. Mousing over the text *other users* displayed a pop-up list of user names (Fig.

1, F). Clicking on the *other users* hyperlink opened a page dedicated to that URL; each bookmark instance and associated metadata were itemized. Only the owner of a bookmark could edit or delete it (or view if private), but all users could copy any public bookmark to their own collection or email it to a colleague or distributed list. Users could pivot browse on any tag (to show all bookmarks tagged with that keyword) or any user name (to show all bookmarks in that user's collection).

Recent popular tags (Fig. 1, B) highlighted the latest topics of interest with a list of the top most frequently used tags over the past seven days. The most popular 30 tags over all time were listed in descending order (Fig. 1, C). Recent tag searches (Fig. 1, D) showed the 10 most frequent searches from the past week; these were displayed with the total counts over all time.

One novel feature was the integration with our corporate directory as a visual promotion of the sense of community. *onomi* displayed photos of users on their own bookmark page and also photos of other users who had either bookmarked the same URLs (*related users by bookmarks* - see Figure 3) or who had used the same tag (*related users by tags*). These features were specifically implemented for our corporate users to help foster social networks. Mousing over any user's photo displayed a pop-up of user details (Figure 3), such as total bookmark count, contact information, and organizational affiliation. Clicking on the bookmark icon next to a user opened the user's bookmark collection.

Users could view all bookmarks they had in common with another user by clicking the + sign next to a related user's name (Figure 3). Alternately, users could navigate the system by manually entering a URL to specify user and tag combinations. For example, *http://onomi/bookmarks.php/**tina**+**paul*** would show all bookmarks that Tina and Paul had in common, and *http://onomi/bookmarks.php/**tina**+**paul**/**css*** would display all bookmarks that Tina and Paul had in common that had also been tagged with the term *css*.

One could also browse all users of the system or browse users by their organizational communities (e.g., department, division, or center). This feature allowed users to explore and discover topics of interest within the company and also to identify organizations that specialized in a particular domain.

Users could search by text string for bookmarks within their own collection, someone else's collection, or the entire *onomi* collection. Search by text included bookmark title, URL, description, and tag. Users could also retrieve bookmark collections specifically by tag, file type (e.g., pdf, doc, xls), user name (i.e., standard user id, first name, or last name), or organization.

In a tag view (i.e., all bookmarks tagged with the same particular keyword), *onomi* listed related tags (i.e., all tags used in conjunction with the specified tag). A user could reduce the subset of displayed bookmarks by creating a filter of related tags. In other words, one could view all bookmarks tagged with both *tagging* and *workshop* (*http://onomi/bookmarks.php/**tagging**+**workshop***) or all of a specific user's bookmarks tagged with both *tagging* and *workshop* (*http://onomi/bookmarks.php/**joe**/**tagging**+**workshop***).

To take advantage of the larger public community knowledge, *onomi* could also display the top 15 bookmarks from *del.icio.us* tagged with the same term. We are currently working on integrating *onomi* tags with corporate services such as internal Google search and other prototypes including enterprise podcasting and a mash-up service. Eventually, *onomi* users would be able to retrieve resources by tag across many corporate applications and services.

Users could add bookmarks either by entering all information manually into a form or by using one of the two browser bookmarklets which enabled click-to-add simplicity when viewing a web page or any document with an associated URL. In addition, users could import all their *del.icio.us* bookmarks or any bookmarks previously stored in a browser. When adding a bookmark via bookmarklet, the title and URL were automatically filled in. Users had the option of entering comments and relevant tags. If a URL had already been bookmarked by another user, *onomi* would recommend previously-used tags. *onomi* also provided type ahead completion, displaying tags from a user's collection that matched what was being typed.

Before the pilot, many of our users already had a bookmark collection on *del.icio.us*. In order to make it easier for our users to maintain both accounts, we enabled simultaneous posting to both social bookmarking services although, for security reasons, we restricted URLs from certain domains from being posted externally. One incompatibility was the tagging syntax; *onomi* allowed multi-word, comma-separated tags, whereas *del.icio.us* used spaces to delimit tags. We resolved this by automatically removing the white space in multi-word tags (i.e., *social bookmarking* to *socialbookmarking*) before posting to *del.icio.us*.

We also integrated bookmarks with email so that users could post a resource to their collection and send out a link to their colleagues at the same time. Associated metadata (i.e., user tags and descriptions) was automatically included with emailed bookmarks. Any bookmark could also be emailed at any time after

it had been posted. The integration with email supported current corporate work practices of emailing resource links to ListServ subscribers with additional benefits: all bookmarked resources sent to a particular distribution list could easily be retrieved later and, because anyone could access public bookmarks on *onomi*, emailed resources would no longer be restricted just to those on the distribution list and could potentially reach a wider audience.

Bookmarks could either be marked public, private, or broken. The individual counts of bookmark types were listed at the top of each user's bookmark page (e.g., *150 public, 5 private, 3 broken*). All public bookmarks were viewable by all users, but private and broken bookmarks were viewable only by owner. (Tags associated exclusively with private bookmarks were also kept private.) A bookmark could only be marked broken by the system although users had the option of overriding the setting. Broken bookmarks were detected via a link scanning service we integrated for housekeeping purposes. Bookmarks with broken links were flagged with a broken link icon, assigned a reserved system tag, and made non-public. The bookmarks remained viewable to the owner, providing users with the choice of fixing the broken link, deleting the invalid bookmark, or doing nothing. Users were periodically notified via email with a link to all their broken bookmarks. Alternatively, users could click on the broken link tally at the top of the bookmark page to retrieve the full list of broken bookmarks.

3. Research questions & evaluation plan

Our general goal was to determine the utility of social bookmarking within our corporate environment. Specifically, our research study questions focused around adoption, usage, and social influences. How would we get our employees to include *onomi* as part of their work practice, and would they continue to use the tool over time? Could social bookmarking tools be useful to our employees, and in which ways: personal bookmarking, information discovery, sharing and disseminating information, finding experts? How would use and behaviors change over time? Would communities develop?

We originally set our pilot plan for six months but extended it to one year due to delays in moving to appropriate hardware and a scalable database. We first launched the prototype to select user communities who were willing to provide us with iterative feedback. Once the prototype stabilized, we introduced it step-wise to other communities through demonstrations and advertisements. Although we had yet to make a formal, corporate-wide announcement, we began monitoring *onomi* adoption by the general audience.

We collected both quantitative and qualitative data on the use of *onomi* and individual users' experiences. Quantitative data was obtained from web server logs and daily statistics derived from database metrics. We used the data to monitor the adoption rate, gain insight into utility through usage patterns, and look for social trends. We also used the data to improve the system during development. Along the way, we received unsolicited feedback from our users on their experiences with the tool, suggestions for improvements, and ideas for new features. In the near future, we plan to distribute questionnaires or surveys for feedback on usage and experiences as well as continue our informal interviews.

We noted several limitations of our initial data collection: the logs did not distinguish between *copy* and *add* events, the number of bookmarks imported was not being captured, *Scuttle* code re-ordered bookmark tags alphabetically so we could not analyze tag order (some researchers believe tag order is significant [10]), and we lacked access to both email logs and proxy server logs so we could not record emailing of bookmarks or user click-throughs on bookmarks.[1]

The next sections present and discuss our open questions and findings to date on each of three areas: adoption of social bookmarking, usage, and practice and social influences.

4. Adoption

Operating in a large, decentralized corporation, we questioned whether our employees would even try a social bookmarking tool. Based on preliminary interviews with random employees, some were already users of *del.icio.us* or other tools, some had never heard of social bookmarking or *del.icio.us*, and some were familiar with the concept but did not see any beneficial applications. Even some of the current users of social bookmarking tools on the Internet saw no need for an internal bookmarking tool that would be inaccessible outside of the corporate intranet.

Our initial problem was in getting people to try the tool. The challenge, however, was in educating people on its use and potential benefits so that they would continue to use it over time. We initially targeted "friendly" corporate user communities, such as librarians and other information stewards, who were

[1] Some of these logging deficiencies had since been corrected, but there remained gaps in our data collection.

willing to help populate the tool with resources relevant to the company as part of their daily research activities. This was an important step for future adoption by others; a resource-barren system would be unlikely to attract potential users of information content or motivate people to use it. We also talked to and encouraged several project teams to start using *onomi*. We began by giving briefings and demonstrations to select groups and, through word-of-mouth, were invited to meet with other communities of interest, organizations, and projects interested in learning more about social bookmarking and the *onomi* prototype.

As new, unsolicited users began to appear on *onomi*, we employed various marketing strategies and began to monitor the adoption rate over time to assess the effectiveness of those strategies (beyond the scope of this paper). First, we produced a series of advertisement banners to appear periodically on the home page of our intranet, designed to be "teasers" which included catchy photographs or gimmicky expressions. We followed the banners with animated announcements on large liquid crystal display screens located in buildings across the company campus. The announcements advertised five informal coffee shop demonstrations which were held over a two week period on one of our campuses. We also distributed our own "schwag": real bookmarks with tassels highlighting *onomi* features and presenting useful tips. Shortly after the demos, we set up a booth at our annual corporate technology symposium (held in two locations), a kind of "science fair" for our sponsors, other contractors, and our own employees.

Another "free" form of advertisement was through integration with other corporate prototypes (e.g., an enterprise podcasting pilot) and interfaces (e.g., shared community spaces and user home pages) either via web services or RSS feeds. These services and sites allowed people to, e.g., post and tag resources directly to *onomi*, retrieve related *onomi* bookmarks, and display a user's popular tags and bookmarks.

Almost 1000 employees had visited *onomi* as of this writing, representing approximately 17.5% of the company, but the challenge of educating potential users to the benefits of *onomi* remained. Of those 1000, only 272 either maintained a bookmark collection or browsed the system on a recurring basis in the past three months. We collected use cases from our more active users and user communities which we made available on the home page for existing and potential users to peruse in the hope that real experiences would be both informative and influential.

5. Usage and statistics

The data presented here represented ten months of user activity and was based on analysis of web server logs and system metrics and supported by interviews with users. At the time of this writing, *onomi* had a total of 7000 bookmarked resources and over 26,000 tags. (Not all bookmarks were tagged, especially those bookmarks imported from a user's browser by batch process.) One fifth of all resources bookmarked were internal to the enterprise (this percentage remained consistent across the ten month window), providing a good case for a closed, corporate social bookmarking system. While *onomi*'s most prolific user had well over 600 bookmarks, the average information provider had 40 bookmarks tagged with 3.7 tags per bookmark, drawn from a set of approximately 29 unique tags.

Although only 175 users had bookmark collections, there were a total of 272 users who had been active in the past three months (not inclusive of users who browsed just once or twice). Overall, the most popular user activities included viewing bookmarks by topic (i.e., tag), perusing one's own bookmark collection, posting bookmarks, and browsing other user collections.

To get more insight into what users were doing and how that activity might change over time, we profiled user behavior by plotting individual browsing activities against bookmarking activities as a percentage of overall user activity on bubble charts in three time windows (Figures 4, 5, 6). Each bubble depicted a single user; the size of the bubble indicated the relative amount of activity over the given time period. The amount of activity reflected a count of each contributing event (e.g., a single bookmark posting, the viewing of a user's bookmark collection, searching on a particular tag, or accessing the help page).

At the lower right of the graphs were the information providers, i.e., users who added, imported, edited or copied public bookmarks (whether intended for their own personal use or to share). In the upper left were the information consumers: those who browsed, searched, and viewed other users' bookmarks and tags. Users at the extreme ends of the diagonal line were either providers or consumers but not both. Those in the middle depicted users who were both providers as well as consumers. Finally, the smattering of smaller bubbles beneath the diagonal line represented users who more often frequented the help page, performed lookups and searches across all data, set up RSS feeds, or just clicked through the navigation bar.

Figure 4 illustrated user activity over a four month window during the initial marketing phase. The graph

depicted 440 users, many of whom were just exploring the system: clicking through the navigation bar, following links randomly, and posting a few bookmarks. Some regular users began to emerge (i.e., larger bubbles), and there was a distinct population of information consumers who were not contributing any resources to the system.

Figure 5 showed the activity of 365 users over the next three months. During this time, there were fewer system "dabblers" and more regular users who posted bookmarks, browsed, and searched. In general, the most active users tended to be information providers.

The data in Figure 6 was collected during the summer months when many employees were on vacation. The seasonal effect resulted in fewer users (272) and less system activity over all, but encouragingly, there was a much higher activity level per user. The extreme information consumers (23% of all users: those who observed without contributing) were becoming even more segregated while the population of users who both added bookmarks and explored other collections grew in size. The percentage of information providers reached 64%.

6. Practices and social influences

Crucial to determining the success of our pilot was to show that social bookmarking was useful to our employees. There were many ways *onomi* could be used: as a personal bookmarking tool, as a mechanism for sharing and disseminating information, for information discovery, to help form social networks, and for finding experts. We hoped to demonstrate the utility of the tool in as many of these areas as possible.

In addition, we were interested in understanding how social aspects might influence the use of *onomi* and what benefits social aspects would provide. Were people more likely to view or copy the most popular bookmarks? Were people more likely to adopt the use of other popular tags or modify their own tagging behavior [11]? Would communities develop? Could *onomi* be used as an expert finding system?

6.1. Personal bookmarking

Users were more likely to examine their own bookmark collection (70%) than other users' bookmark collections and just as frequently looked at their own tag sets versus other people's tag sets.

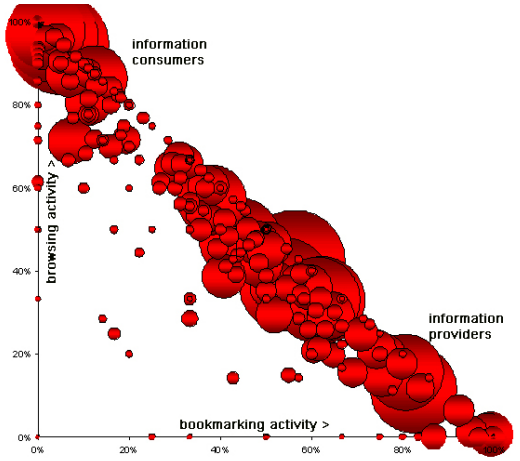

Figure 4 Snapshot of activity: Nov 8, 2005 – Mar 7, 2006

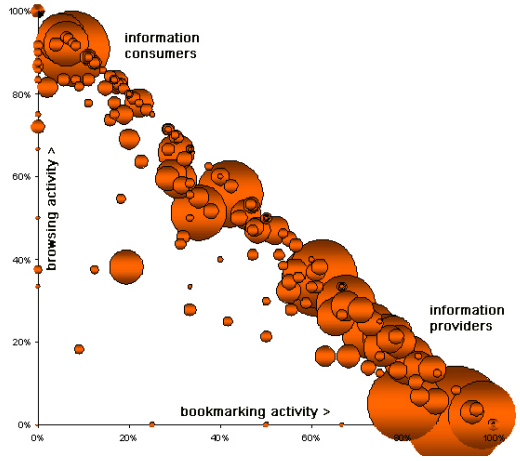

Figure 5 Snapshot of activity: Mar 8 – Jun 7, 2006

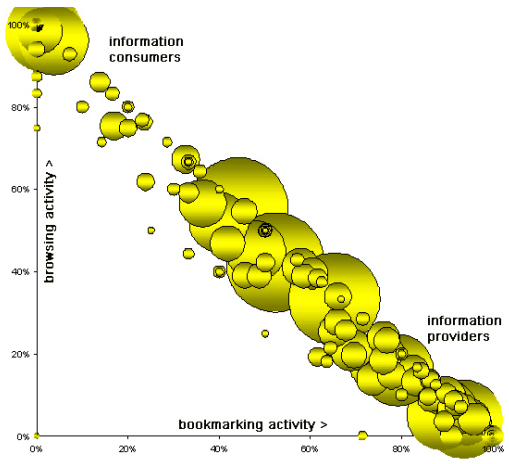

Figure 6 Snapshot activity: Jun 8 – Sept 7, 2006

Private, enterprise-internal bookmarks, accounted for 2% of all bookmarks. Based on personal experience, however, the authors presumed that these private bookmarks might have linked to performance reviews, documents stored locally, and human resource materials – personal resources that users might otherwise have bookmarked in a browser. A much higher number of private bookmarks (9%) came from the Internet. We noted that users who imported their bookmarks, either from a browser (13% of users with bookmark collections) or from a *del.icio.us* account (8%), did a batch import, often marking all bookmarks private and then later tagging each of those bookmarks and making some of them public. This practice may have falsely inflated the number of private, external bookmarks but it did show that users were transferring personal resource management to *onomi*.

Publicly, users did not only post work-related resources but also links to personal interests, sports and hobbies, and photographs as well as pointers to favorite restaurants, venues, and local events. One user posted bookmarks specifically for his own use but did not care to make them private: "*I generally bookmark things I want to remember and read again. I do not bookmark items with the intention of sharing (although I do not make them private).*"

6.2. Information sharing & dissemination

89% of all bookmarks posted were public. (Note that bookmarks were marked public by default.) Information providers at the far lower right of Figures 4-6 could consist of both users who bookmarked just for themselves as well as "mavens" [12] who posted for the sole purpose of assisting others (e.g., librarians). One information provider stated, "*I use onomi not only as a way to save interesting sites and articles that I have come across, but I use it to promote awareness of interesting articles for others, and as a way to publicize the availability of new … resources*."

While we lacked exact statistics on emailed bookmarks, there was some evidence (e.g., publicly archived email) that users employed this feature to distribute resources to individuals, project teams, and ListServ communities.

6% of our users created *onomi* RSS feeds to shared community spaces, employee web pages, and project sites. This practice, together with integrated feeds to other corporate prototypes, was helping to distribute information beyond just the *onomi* user population.

We interviewed one particular consumer, a project lead, who told us that his project members decided to bookmark papers, works-in-progress, deliverables, and resources relevant to their research in order to facilitate sharing. They created their own project-specific tag to use, making retrieval straightforward. Although the project lead had no bookmarks of his own, he accessed *onomi* for the sole purpose of keeping up to date with his team's work. He explained that "*the value of onomi is disproportionate to my amount of use.*"

6.3. Information discovery

The user profiling in Figures 4-6 was inspired by one user who described himself as a "leech;" he was not interested in sharing his own resources but often explored other users' collections or bookmarks tagged by a specific topic to discover new resources. This particular person had also created a number of *onomi* RSS feeds. This type of user would fall at the extreme upper left of the graph, as an information consumer like the project lead described above.

"Folksonomies are predisposed to discovering unknown and unexpected resources" [13], and our usage data supported this. When looking at bookmarks by topic, users were more likely to look at all bookmarks by tag (75%) than their own (22%) or a particular user's (3%). This suggested that people were interested in discovering new resources contributed by the larger population.

Because we did not have bookmark *copy* metrics available until the last few months of this study, we could not always distinguish between a bookmark copied from another user's collection and a resource found independently. However, we could infer from duplicated tags that some bookmarks might have been copied, and there was later logged evidence that at least 11% of users with bookmarks were in the habit of copying from others. Of the 7000 bookmarks, 8% were shared with more than one user. Although one might expect this percentage to increase over time, some users stated that they had no need to copy someone else's bookmark since they could always find it easily. One user said that she would only copy another user's bookmark if she wanted to add different tags that would help her remember how to find the bookmark later [14]. (Click-through data would have greatly benefited our analysis here.)

The most popular bookmarks, highlighted on the *pop-onomi* page, were shared by no more than seven users. Of interest was an observation that bookmarks on both the home page (i.e., recent postings) and the popular page were more likely to be bookmarked by others than bookmarks buried in user collections [13]. As the number of people bookmarking the same item increased, so might its likelihood of being copied by yet other users. One user based his trust on this factor:

"*I rely on a 'quality of link' predictor; I look at how many people have bookmarked certain things.*"

6.4. Forming communities and social networks

Some project teams extended their communities virtually by agreeing on specialized tags for their bookmarked resources, as described above. There were indications that other virtual communities began to form as tags began to converge. The number of shared tags had increased by 140% in the past three months, and users stated that they changed their tags or added new tags based on what others with similar interests were using. Indeed, the average number of tags per bookmark increased from 2.6 to 3.7 during the study.

One user sent glowing feedback to the development team after the release of our *related users* feature; she was now able to see others who were interested in the same topics. More and more users were discovering the *related user filter* (viewing shared bookmarks or shared topics) which we believed would help support social networks. In the next prototype release, users would have the ability to comment on bookmarks. Comments would be associated with the URL and thus accessible through any bookmark of that URL. We would be interested to see how this new feature might be used, and if it would help establish new communities and support existing ones.

6.5. Expertise finding

Although users more often viewed all system bookmarks by topic, there were some cases of people looking at a specific user's bookmarks by topic. This might have been an indication of identifying others as potential experts on a particular topic. Users themselves had begun to feed the system with data to support expertise finding. The authors observed one user who had been rapidly bookmarking select employee directory pages and tagging them with the term *expert* along with other tags referring to a skill or research area. As more and more of these *expert*-tagged bookmarks began to appear on the home page as recent postings, the activity must have caught another user's attention; he bookmarked his own directory page with the tags *expert* and *design*.

While *onomi* alone might not suffice as an expert finder, we believe it would complement or feed corporate expert finding mechanisms. For example, users' most used tags might indicate expertise in a particular area. As the system slowly approached critical mass, these and other statistics would provide further insight.

7. Insights and potential benefits

One critic claimed that tagging systems, such as social bookmarking tools, are "simple enough for people actually to use and robust enough to be of value to a community of users, not just the one doing the tagging" [15]. Social software tools on the Internet, such as *del.icio.us* and *flickr* (both acquired by Yahoo in 2005), had shown that the use of tags could be effective in information discovery and retrieval even with only a small percentage of users employing them. Caterina Fake, cofounder of *flickr*, reported that a community could benefit when just 15% of its members were tagging information [16].

We believed that bringing social bookmarking inside a corporate environment would provide even greater benefits both to individual users and the company as a whole. Despite a small user base, resources bookmarked in *onomi* reached many employees and corporate sites. We showed that users were importing bookmark collections and sharing them with others. Bookmark and tag feeds populated internal websites, shared spaces, and other corporate applications. Users bookmarked internal resources at a constant rate, with the assurance that corporate knowledge was protected from the outside world. Virtual communities began to develop in *onomi* and specialized tags emerged around those communities. Project teams and groups used social bookmarking to share and distribute resources.

onomi replaced some existing tools and work practices while supplementing others. In some cases, it replaced the use of browser bookmarks and possibly the use of external social bookmarking tools (or supplementing them) by allowing users to share and protect their internal resources. While social bookmarking could not replace commonly used repositories, it helped to make the use of multiple repositories transparent (corporate repositories and collaborative spaces were scattered across the intranet). Similarly, it saved the librarians from having to re-gather multiple resources for different customers; they could tag resources appropriately when they first identified them and then later retrieve them. These bookmarked resources then became discoverable by others. As the system became populated with more and more relevant resources, it saved users from having to do their own exhaustive searches; the fact that someone else had found and bookmarked a particular resource was a rating, in a sense. Lastly, *onomi* augmented corporate knowledge distribution through

RSS feeds as well as through ListServ (by enabling simultaneous bookmarking and emailing).

In our study, the benefits of social bookmarking extended far beyond simple resource management; the tool was valuable in a corporate environment for sharing and disseminating, promoting information discovery, supporting communities and social networks, and feeding expertise finding.

8. Issues and future challenges

As we move forward with what appears to be a successful pilot, we face numerous issues and challenges. Many users and corporate groups have requested both group-based (e.g., teams, projects, organizations, communities of interest) and role-based (e.g., librarian) accounts. Because of the single sign on implementation, based on the directory server, *onomi* currently supports only standard user-based accounts. While the use of specialized tags helped create virtual communities for some groups, this workaround does not support group access to private or restricted bookmarks. We will investigate alternate methods in supporting different types of user accounts.

Another issue we face is the handling of bookmarks and tags when an employee leaves the company. A solution to the previous issue may also lend itself to this problem of stewardship. Re-assigning valuable resources to an "ex-employee" account is preferable to removing them.

We are investigating new directions for tagging. We are considering incorporating semantics (e.g., is *java* a programming language or a coffee?) or the ability to type and sub-type a tag (e.g., person: *john*, location: *office*). Changing our tagging infrastructure will require considerable collaborative effort with developers of other corporate services and applications that either are currently or will be integrated with *onomi*. Ideally, we would like to use the tags from all systems to help feed our corporate subject taxonomy, and there is support from our users: "*Tags are surprisingly useful considering we are relying on an unstructured mechanism. I am really surprised how accurate tags are and how precisely they are used*." We would also like to improve tag recommenders and predictors, perhaps using machine learning techniques or information extraction.

As *onomi* continues to expand its user base and resources, we will continue to monitor the data not only as feedback for improving the prototype but also to provide further insight into the utility of social bookmarking in a closed community.

10. Acknowledgments

The authors would like to acknowledge Dave Lehman (CIO during *onomi* conception and development) for his support, funding, and championship of the social software and semantic web work at MITRE.

11. References

[1] *del.icio.us* http://del.icio.us/

[2] *Connotea* http://www.connotea.org/

[3] Green, Heather. Tagging: Keeping Tabs on the Net in *BusinessWeek*, September 9, 2005.

[4] IBM's *dogear* http://domino.research.ibm.com/cambridge/research.nsf/0/1c181ee5fbcf59fb852570fc0052ad75?OpenDocument

[5] Millen, D., Feinberg, J., Kerr, B. Social Bookmarking in the Enterprise from *Social Computing* Vol 3, No. 9, November 2005.

[6] Rox, T. Social Bookmarks Charts, version 3.0, Jan 2006. http://www.irox.de/roxomatic/856

[7] Hammond, T., Hannay, T., Lund, B., Scott, J. Social Bookmarking Tools (I), A General Review in *D-Lib Magazine*, Volume II, Number 4, April 2005.

[8] S*cuttle* http://scuttle.org/

[9] Todras-Whitehill, E. 'Folksonomy' Carries Classified Beyond SWF and 'For Sale', the New York Times, Oct. 5, 2005.

[10] Golder, S., Huberman, B. Usage Patterns of Collaborative Tagging Systems in *Journal of Information Science*, 32(2). 198-208. 2006.

[11] Porter, K. Folksonomy Notes: Considering the Downsides, Behavioral Trends, and Adaptation in *Bokardo, A Blog About Web 2.0* http://bokardo.com/archives/folksonomy_notes

[12] Gladwell, M. The Tipping Point: How Little Things Can Make a Big Difference. Publishers: Little, Brown, 2000.

[13] Kroski, E. The Hive Mind: Folksonomies and User-Based Tagging, http://infotangle.blogsome.com/2005/12/07/the-hive-mind-folksonomies-and-user-based-tagging, 2005.

[14] Biddulph, M. Introducing Del.icio.us. XML.com. http://www.xml.xom/pub/a/2004/11/10/delicious.html, 2004.

[15] Folksonomy Schmolksonomy in PeerPressure, the Official AllPeers Blog. http://www.allpeers.com/blog/index.php?p=36, 2005

[16] Fitzgerald, M. The Name Game in *CIO Magazine* http://www.cio.com/archive/040106/et_main.html, 2006.