

Dogear: Social Bookmarking in the Enterprise

David R Millen, Jonathan Feinberg, Bernard Kerr

One Rogers Street. Cambridge, MA 02142

{david_r_millen, jdf, bernard_kerr}@us.ibm.com

+1-617-693-7490

ABSTRACT

In this paper we describe a social bookmarking service designed for a large enterprise. We discuss design principles addressing online identity, privacy, information discovery (including search and pivot browsing), and service extensibility based on a web-friendly architectural style. In addition we describe the key design features of our implementation. We provide the results of an eight week field trial of this enterprise social bookmarking service, including a description of user activities, based on log file analysis. We share the results of a user survey focused on the benefits of the service. The feedback from the user trial, comprising survey results, log file analysis and informal communications, is quite positive and suggests several promising enhancements to the service. Finally, we discuss potential extension and integration of social bookmarking services with other corporate collaborative applications.

Author Keywords

social software, social bookmarking, tags, folksonomies

ACM Classification Keywords

H.5.3 Group and Organization Interfaces: Collaborative computing.

INTRODUCTION

One of the greatest challenges facing people who use large information spaces is to remember and retrieve items that they have previously found and thought to be interesting. One approach to this problem is to allow individuals to save particular “search” strings, in order to recreate the search in the future [2]. Another approach has been to allow people to create personal “collections” of material, such as the electronic citation bundles (called binders) in the ACM Digital Library. Collections of citations can be created manually by readers or through execution (and alerting) of a saved search.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2006, April 22-27, 2006, Montréal, Québec, Canada.

Copyright 2006 ACM 1-59593-178-3/06/0004...\$5.00.

Perhaps the most familiar approach to “re-finding” information on the web has been the use of personal bookmarks, as supported by various web browsers. For example, the Mozilla browser supports the creation of collections of URLs, which can be annotated using keywords or free-form text, and then sorted on a variety of dimensions (e.g., time last visited, keyword, location). An early study of bookmark use showed that people created bookmarks based on the quality of and personal interest in the content, high frequency of current use, and a sense of potential for future use [1]. Furthermore, the number of bookmarks contained in an individual collection grew steadily and roughly linearly, and the use of folders to categorize bookmarks increased as the size of the collection increased. A single level of folders was reported for collections with fewer than 300 bookmarks, while a multi-tiered hierarchy was used for larger collections.

The desire to find and share information among small groups, teams and communities of practice has led, not surprisingly, to the development of a number of shared bookmarking systems. Early shared bookmarking systems often used automated techniques to support the creation and categorization of collections of web bookmarks [8,9,16]. These innovative systems met with some success, although they consistently seemed to fall short of their potential use. Several explanations for their limited success have been offered, including limited privacy protection, little support for end-user tailorability and high requirements for active user participation [7].

SOCIAL BOOKMARKING SYSTEMS

Recently, there has been an emergence of shared bookmarking applications. Their tremendous popularity and growth of use have prompted a second look at this kind of collaborative software. Introduced in 2003, the del.icio.us [3] social bookmark manager was one of the first of this kind of application, and has enjoyed an early and large base of committed users. A flurry of similar offerings have since been unveiled, including a recent beta release of Yahoo’s “My Web 2.0” [11].

These systems share a number of features. First, these social bookmarking tools allow individuals to create personal collections of bookmarks and instantly share their bookmarks with others. These centrally stored bookmark collections bring immediate personal benefit by providing a col-

lection that can be browsed from any web-accessible machine.

A second, and significant, enhancement in these systems is the use of keywords or tags that are explicitly entered by the user for each bookmark. These tags allow the individual user to organize and display their collection with labels that are meaningful to them. Furthermore, multiple tags allow bookmarks to belong to more than one category, avoiding one of the limitations of the hierarchically organized folders found in most web browsers. The use of tags to create an emergent classification system has been controversial and is likely to spawn significant research in the short term [15].

The final distinguishing characteristic of these social bookmark applications is the social nature of their use. There is a bias towards increased transparency in these tools. While bookmark collections are personally created and maintained, they are typically also visible to others. A number of user interface elements allow social browsing of the bookmark space. For example, user names are “clickable” links, and, when a name is clicked, the bookmark collection for that user is presented. This allows someone to get a sense of the topics of interest for a person. Tags are also clickable, and when selected will result in a list of all bookmarks that share that tag. This is a useful way to browse through the entire bookmark collection to see other information sources of interest. We call this ability to reorient the view by clicking on tags or user names, “pivot browsing”; it provides a lightweight mechanism to navigate the aggregated bookmark collection.

SOCIAL BOOKMARKING FOR THE ENTERPRISE

While internet social bookmarking services, like del.icio.us, are of great benefit for publicly accessible web resources, they are problematic for web resources on most corporate intranets. While services like del.icio.us do allow bookmarking of intranet pages, corporate firewalls prevent access to these resources to anyone outside the organization. Furthermore, public sharing of bookmarks to intranet resources may be of concern as proprietary information, albeit limited in nature, could be leaked.

Nevertheless, the apparent success of internet-based social bookmarking services begs the question of whether large enterprises or organizations would also benefit from a social bookmarking system. We have been investigating this question by designing and deploying an enterprise-scale social bookmarking system. In this paper we describe the important design decisions we made for the “dogear social bookmarking service” (henceforth referred to as *dogear*). We will then provide a description of the primary characteristics of *dogear*, including a detailed presentation of the user interface. And finally, we provide the results of a field study of the *dogear* service.

DESIGN PRINCIPLES

Identity and transparency

The first significant design decision was whether user “identity” in the application would be based on real names, or whether pseudonyms would be allowed and supported. We decided to require the use of real-world identity in the bookmarking application for the following reasons. First, one of the expected benefits of the system was to allow users to make inferences about the interests and expertise of others based on informal browsing of bookmark collections. This discovery of people with shared (or complimentary) interests would help to nurture the communities of practice within the enterprise, and potentially allow searching for and finding experts on specific topics to help solve business problems. Real name identity would also allow users to look up additional information about other users in various corporate databases (e.g., corporate online directory and enterprise web). In addition, real name identity would facilitate communication between users of the application since the various corporate collaboration tools (e.g., corporate directories, email, chat,) all use real name identities. Finally, there is a strong cultural norm within the organization to use more formal names within the corporate applications.

It should be noted that while creating personal bookmark collections requires authentication of real name identity, it is possible to browse the enterprise bookmark collection anonymously. This invisible participation is a form of “lurking” and allows the benefits of bookmark sharing to extend to the entire enterprise. Lurker behavior is counted in application statistics (e.g., number of times a bookmark is “clicked” to access a particular information resource). While lurking in online environments has been characterized as a form of “free riding,” it has been shown to provide an important supportive role in the sociability of online environments [12].

A second important design decision concerned the degree of support for access control within the application. While the early public internet bookmarking services (e.g., del.icio.us) have provided open access to everyone’s bookmarks, there are valid reasons to consider limiting access to some set of bookmarks. For example, a group may want to create a collection of information sources in support of a project that may be extremely confidential or bound by contract to restricted access. Teams may want to create shared bookmarks that are visible only to the group. We could also see a need for “role” or job-based collection of information sources. For example, managers might benefit from bookmarking and sharing certain information sources, while the sales force would benefit from another. And finally, we saw a need for private bookmarks. An informal survey of potential users of the system indicated that many people had bookmarks that they considered personal and would not want to make public. We did not want to force people to use multiple bookmarking tools—a personal bookmarking tool (e.g. in a browser) and a corporate tool for shared

sources. Our initial implementation includes both private and enterprise “public” bookmarks. We have architected solutions to both team and role based collections of bookmarks and will roll these capabilities out at some time in the future.

The decisions to use real name identity and to support private bookmarks were made with some trepidation. The success of social software applications, in general, requires “critical mass” of user participation in order to provide value to users and to ensure sustainable contribution levels and vibrant interactions. Real name identity may discourage some people from using the system, while private bookmarks will significantly reduce the benefits of information sharing among users. For a large enterprise, we believe that the ability to reach critical mass will not present a problem.

DOGEAR SOCIAL BOOKMAKING SERVICE

The *dogear* service prototype was designed to simply and elegantly display each bookmark’s content within a naviga-

tion model that allows users to manage and explore the collection in different ways. The user’s bookmark collection is a reverse chronological list of their most recent bookmarks, similar in format to a blog, as shown in Figure 1. Each bookmark has a number of pieces of metadata which give the user useful information about its context and content. Other pages in the system include lists of recent, popular, and tag collections, along with specific URL and search result views. The following is a detailed description of the how we expose the metadata in the interface, the primary elements of the navigation, and some of the other key features of the system.

Each bookmark has a title and description (A). By default the title is taken from the HTML “title” tag of the bookmarked page when the bookmark is made, but can be edited. The description is optional but we offer a suggestion from the URL’s “description” tag if there is one. Whenever possible we also add a “favicon” to each bookmark as it is created.

The screenshot shows the user interface for Jack Russell's bookmarks. The header includes the user's name, navigation links (All, Popular, My Bookmarks, My Subscriptions), and utility links (Settings, Help, Feedback, New). A search bar and a dropdown menu for 'All Bookmarks' are present. The main content area displays a list of bookmarks, each with a title, date, author, and tags. Annotations A through K point to various UI elements: A (bookmark title), B (author), C (date), D (comments), E (copy link), F (tags), G (description), H (user profile), I (dogear logo), J (tag filter), and K (tag list).

Bookmark Title	Date	Author	Comments	Copy	Tags
Jeffrey Zeldman Presents The Daily Report	15 SEP 2004	JACK RUSSELL	3	COPY	blog, css, design
Stylegala ??? the finest CSS and web standards resource	10 SEP 2005	JACK RUSSELL	0	COPY	css, design
Simon St. Laurent	1 SEP 2005	JACK RUSSELL	13	COPY	blog, css, design
Curvelicious	22 AUG 2005	JACK RUSSELL	0	COPY	css
Intranet CSS Standards Guide	02 APR 2005	JACK RUSSELL	0	COPY	blog, css, design
No More Tables CSS Layout Techniques: What brings us here today...	02 JAN 2005	JACK RUSSELL	0	COPY	css, design, tutorial

At the bottom, there is a pagination control showing 'Show 10, 15, 25, 50 items per page. Showing 6.' and a footer with 'RSS ATOM JAVASCRIPT'.

Figure 1. Illustrative screenshot of the dogear social bookmark service. This view is of an individual’s collection of bookmarks, showing only the bookmarks with “design” and “web” tags.

Other metadata include the time the bookmark was made (B), the author (C) and any comments made on that bookmark (D).

The user can directly pivot-browse to an author’s bookmarks by clicking on the author’s name, or add comments via the comments link. The comments are identical in style to the comments one can create in response to blog entries. The highlighted numbers (E) indicates the number of other people who have also bookmarked that URL; clicking on it changes the view to a list of bookmarks for that URL made by other users in the *dogear* system.

Two other pieces of metadata which we believe are critical to an enterprise bookmarking system are intranet and authentication identification. These are shown on any bookmark which references an intranet page or requires authentication (F). This clearly indicates to the user any internal page which is potentially confidential or a page which will require a password to access.

Finally, and perhaps most importantly, we dedicate a column to the display of each entry’s user-defined tags. This allows users to clearly see the tag content for each bookmark and, at a glance, to see the number of tags for each entry. Each of the tags is clickable, providing an opportunity to pivot-browse to new views of that author’s collection of bookmarks for that tag.

Beside the bookmark entries we surface interface elements that expose useful data about the users of the *dogear* service, which helps users exploit *dogear*’s enterprise nature. For example, when looking at an individual’s bookmarks a user can learn more about the author via a disclosure triangle beside the author’s name (H). This reveals the author’s job title, location and a photo which are all taken from our internal intranet personal pages. In addition we allow users to hyperlink directly to the author’s personal page, or their internal blog, or begin an email. These features encourage networking and help with expertise location based on an individual’s bookmarks.

Another feature that exploits our awareness of identity is the “People” tab. This contextual tab shows a list of people related to the currently viewed bookmarks. For example, as shown in Figure 2 below, when the user looks at an individual’s bookmarks page the “People” tab contains a list of users who have also bookmarked the same URLs for themselves. This is, in effect, a simple social network of similar people, and pivot browsing on one of their names allows the user to begin investigating their bookmarks. The “People” tab on the latest bookmarks page lists the people who have recently been most active in creating bookmarks on *dogear*.

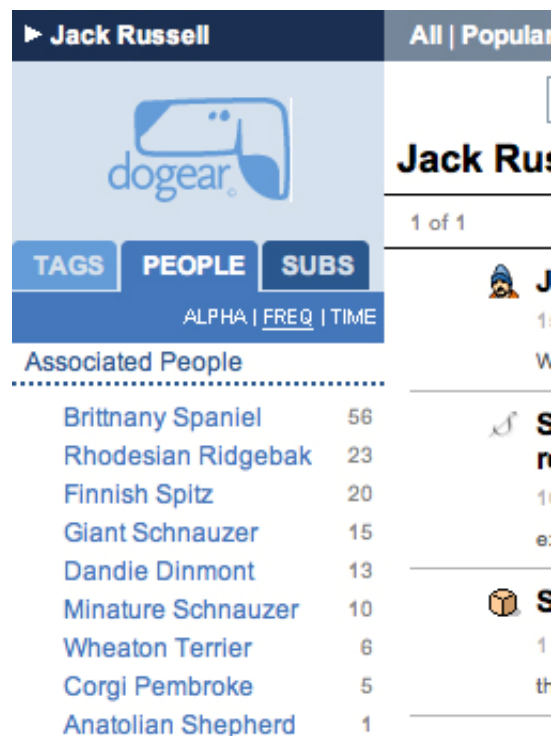


Figure 2. The “people” tab showing a list of *dogear* users who are associated with the current list of bookmarks.

Navigation

One of the challenges in the design of the system was to help the user navigate the many views possible in *dogear*. In particular it was a challenge to show users where they were when examining an individual’s tags. As an illustration, clicking on one user’s tag gives a list of all their bookmarks which have that tag. This can often result in a very long list of bookmarks but the user can reduce this list by specifying another tag that must be used in conjunction with the first. For instance a user might first want to look at someone’s “design” tag, and then realize they need to be more specific. The user might add the “web” tag as another constraint, so that the result is all bookmarks with both “design” and “web” as tags. These intersecting tags are shown as the “Associated” tags list in the interface (J) as seen in Figure 1. As tags are added to the current constraint we represent them directly under the “Tags” tab area (I) and in the page title. This version of pivot browsing is similar to the linear “bread crumb” type navigation interface commonly used in websites, but here the navigation is less linear as the user can also remove any tag from the current view’s requirements. In particular, clicking on the minus-sign icon for a tag removes that tag from the current constraint. Alternatively, clicking on the “All” link will pivot-browse the user to a new view of all of the *dogear* users who have used that tag.

Another navigation technique used in *dogear* is a Tag Cloud Index. Tag clouds have recently become a popular way to represent and navigate tags in web applications

(e.g., Flickr [5], Technorati [13]). In a tag cloud, tags are listed, not in a single column, but rather one after the other in a large paragraph. Each tag is also typographically styled with size and/or color in relation to the frequency that they are used in the system. This has the effect of giving the user a tag list which simultaneously shows frequency and alphabetical information at once. The *dogear* service also provides a tag cloud view of “Active” tags (K).

The *dogear* tag cloud has two additional features. The first is a slider which allows the user to dynamically reduce the tag cloud to the more frequently used tags in the current view. The second feature is a letter index down the left hand side of the cloud, which makes it much easier to scan the cloud to find specific tags.

All navigation actions taken by the user are expressed, ultimately, in an URL-encoded state, as described in the section below, “Dogear: An Example of REST”.

Searching the *dogear* Web Site

The search box in *dogear* works as expected; the user types in a free text query, and a list of bookmarks is returned, in order of relevance. The drop down options allow the user to search by user name (“People”), by tag (“All Tags”), by all fields within the user’s bookmarks (“My Bookmarks”), or within the currently displayed set of bookmarks (“These Bookmarks”). For example, having restricted the currently visible bookmarks to Leon Berg’s “java” tag, one may search within that set for the occurrence of some arbitrary word or phrase in the description or title.

However, because *dogear* uses the Lucene search API, it is possible for sophisticated users to construct rather refined text queries. For example, the following are all valid *dogear* search queries:

“title: sales”

Bookmarks having the word **sales** in the title field.

“tag: wiki AND NOT description: wiki”

Bookmarks bearing the tag **wiki**, but not having the word **wiki** in the description.

“host: microsoft AND description: activit*”

Bookmarks with **microsoft** in the host name, with some reference to **activity** or **activities** in the text description.

“+person: david +collaboration –software”

Bookmarks created by anyone named **david**, which mention **collaboration**, but not **software**, anywhere in the text description, title, or tags.

Although *dogear* does not currently provide a way to explicitly save a search query, one may certainly subscribe to a given query by copying the appropriate URL (RSS, Atom, etc.) from the query result page. We have occasionally observed users bookmarking a *dogear* page on *dogear*, which is an ingenious answer to the need to save state.

Alerting and Discovery

In order to encourage ongoing use of a social bookmarking application, it is important to consider mechanisms to alert users of relevant new information. Indeed, various notification mechanisms have been shown to contribute significantly to sustaining online interaction in other collaboration application [10]. We envision several ways to keep users informed of significant new content.

First, the system currently provides RSS and Atom presentations of every page. This allows users to subscribe to general bookmark collections, such as the most active or most recent bookmarks. In addition, users may subscribe to bookmark collections belonging to a particular individual, or to a particular tag. A standard RSS reader can then be used to monitor new information in a particular collection of bookmarks.

Second, we plan to exploit collaborative filtering techniques to screen new bookmarks for those that are predictably of interest to an individual (or group of individuals). Common interests will be inferred based on a number of observable user actions, including use of similar tags and/or tag combinations, similarity of bookmark (URL) collections, and click streams that indicate interest in specific kinds of bookmarks. Text analysis of bookmark titles, descriptions, and comments will also be used to determine bookmark relatedness. Email notification, system tray status indicators, or small pop-up window alerts can be used to inform users of new bookmarks and other personally relevant information.

DESIGNING FOR EXTENSIBILITY: ENTERPRISE REMIXING

In designing *dogear*, we have considered both the web site user and the more technically sophisticated user who wishes to exploit *dogear* with computer programs. The needs of both can be addressed simultaneously by adhering as closely as is practical to the architectural style that Roy Fielding dubbed Representational State Transfer, or REST [4]. For the user, the REST style is characterized by transparent, meaningful URLs. In turn, the URL structure suggests an implementation that has proven to be flexible, maintainable, and responsive. Our ability to play nicely with other research and production systems on the intranet has led to early adoption by a number of people and organizations within our company..

Dogear: An Example of REST

The REST architectural style specifies that a given resource be represented by a stable URI, such that the same URI will *mean* the same thing over time (and not necessarily that it will yield the same data over time). This implies that the current state of an application must depend entirely on the URI, and not on opaque, hidden conditions, such as cookies. The *dogear* application state consists of a set of constraints on which bookmarks are currently visible, an output format, a current page number, and the number of links per page, as shown in Figure 3

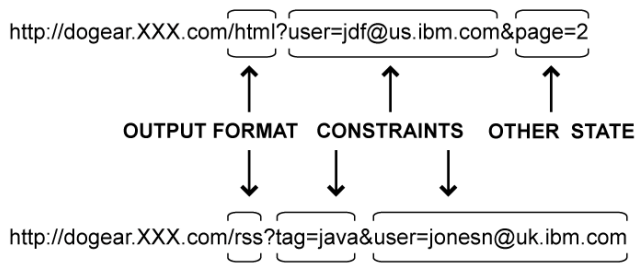


Figure 3. Examples of how *dogear* encodes application state in URLs

Output Format

The first element of a *dogear* URL path selects the desired output format. Currently implemented formats include:

/html—the web site itself

/rss—RSS 2.0

/atom—Atom 1.0

/js—a JavaScript snippet, which, when evaluated, causes HTML to be emitted on the including page

/xbel—the XML Bookmark Exchange Language

Each of the machine-readable formats is extended, via XML namespaces, to include *dogear*-specific metadata.

Programmatic read/write access to ones own bookmarks is provided by an HTTP API, which resides at **/api**. Bookmarks are created and modified via HTTP POST, using simple name=value query parameters to provide bookmark title, url, and tags. Bookmarks may be deleted by calling **/api?url=url** with HTTP DELETE.

Query Constraints

Results may be constrained through the use of HTTP GET query parameters. Any combination of a person, a tag constraint, and a text search may be specified. The person constraint takes the form of **user=userid**, and restricts the result set to those bookmarks created by the specified person. There are three tag constraints: tag, any, and all. The query **tag=foo** restricts results to those bookmarks tagged with foo; **any=foo+bar+baz** will accept bookmarks bearing any of the foo, bar or baz tags; **all=foo+bar+baz** permits only those bookmarks bearing all three of the specified tags. Bookmarks are returned in reverse chronological order, except in the presence of a **searchText** parameter, which causes results to be ordered by search relevance. In the absence of any constraints, the ten most recent bookmarks are returned.

Other State

All queries may include either or both of the **page** and **ps** parameters, which specify the result page and page size, respectively.

These URLs are constructed, and their output parsed and transformed, easily enough to permit very rapid creation and deployment of useful services which retrieve, aggregate, and represent our data with a minimum of programming.

FIELD STUDY

The *dogear* service prototype was launched for a small friendly trial in March, 2005. A small group of committed users quickly formed. A more formal field trial of the *dogear* service within the corporate intranet began in July as part of a new technology incubation program, and announcements about the service availability have been made on intranet web sites. Additional buzz about the availability of the *dogear* service has been created through word of mouth and significant discussion of the service on the corporate intranet blogs.

We have begun to assess the usage patterns of *dogear*'s early adopters based on an analysis of the server log files. Included in the log files are user actions (e.g., create, delete, edit a bookmark, bookmark "clicks"), user and bookmark owner identifiers, and a time and date stamp. In addition, we have analyzed aspects of the bookmark collection itself, which provides additional information about the composition of bookmarks (e.g., tag information).

The user activity analysis presented here was based on log files covering an eight week period during July to September, 2005. During this usage period, 686 individuals were recording using the *dogear* service, with 27% creating bookmarks and 54% clicking on a link (URL) to view a web page that was bookmarked. An analysis of *dogear* users who created at least one bookmark showed that there was some geographical diversity (70% Americas, 7% from Asia-Pacific, and 23% Europe). All were working in a large firm (i.e. more than 300K employees) in the Information Technology industry, with a majority of users involved in software development and support.

A summary of the total bookmark collection can be seen in Table 1. It should be noted that almost one quarter of all *dogear* bookmarks are from corporate *intranet* pages, which supports the design goal of providing a social bookmarking service for enterprise (secure) web pages. Also of note is that the vast majority of *dogear* bookmarks are shared. Less than three percent of all bookmarks are marked private.

	n	%
Total bookmarks	13174	100
Shared Internet	9960	76
Private internet	131	01
Shared intranet	2965	23
Private intranet	118	01

Table 1. Summary of total bookmark collection.

Some usage results from the eight weeks of the field trial can be seen in Figure 4 below. The general pattern shows a sustained number of bookmark additions (CREATES) with a slight increase in use over time. Also visible is evidence that people are following the bookmarks to the information source (shown as “CLICKS” in Figure 4). We see very little bookmark “editing” during this time, which would be evidence of bookmark management activities such as refinement of tags or keywords, or additional annotation to the bookmark. This is not surprising as we are early in the trial and others have shown that bookmark management increases as the size of collections increases [1].




Figure 4. Activity trends on dogear over eight weeks field trial.

Use of tags (bookmark keywords)

Since the use of tags (or keywords) is a defining characteristic of social bookmarking services, we have examined the use of tags in *dogear*. 1971 distinct tags have been used within *dogear*, with a median of 10 distinct tags per user. We observed a strong correlation between the number of tags used and the number of bookmarks created $r = .56$, ($p < .001$). This correlation was stronger for *dogear* users with smaller bookmark collection (< 10 bookmarks), $r = .74$ ($p < .001$) than those with larger collection (> 10 bookmarks), $r = .5$ ($p < .001$). These correlations are stronger than those reported from an analysis of del.icio.us user [6].

With a few exceptions, the number of tags used per bookmark was quite small. The average number of tags per bookmark was 2.3, while both the median and modal number of tags was 1. A summary distribution of the number of tags per bookmark can be seen in Figure 5. The cumulative frequency reveals that over 80% of the bookmarks have three of fewer tags.

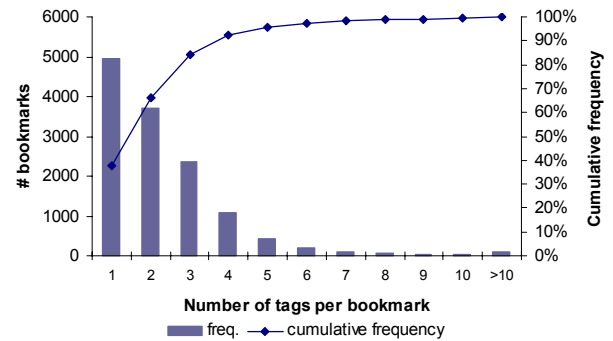


Figure 5. The total number of tags per bookmark, and the cumulative frequency of bookmarks.

We have also begun to look for patterns of shared interest in various information sources. We have used social network analytical methods [14] to begin to understand the information affinities among *dogear* users. A two-mode sociogram showing which individuals have used similar tags for various programming languages can be found in Figure 6. We plan to more systematically mine this information to automatically generate tag groupings and to alert *dogear* users of new bookmarks that might be of interest to them. We have begun to consider the ways we can use this kind of information in helping to share information among larger groups of people within the organization.

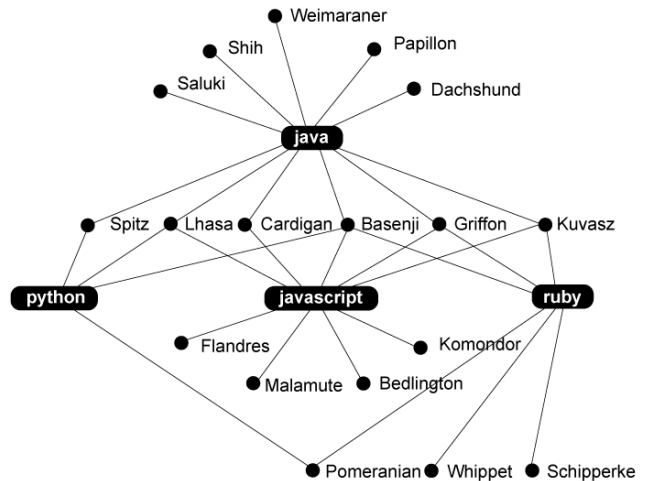


Figure 6. Sociogram showing which individuals have used similar tags (java, javascript, python and ruby).

Navigating through dogear

The analysis of logged user activity is also revealing about the manner in which individuals explore the *dogear* information space. As described above, we have provided the ability to pivot-browse the *dogear* bookmark collection by selecting from any of a number of active links on a page. In Table 2, we provide a summary of the user pivot behavior during the field trial.

	n	%
Self	303	44%
Own tags	86	13%
Other people	235	34%
Specific person's tags	104	16%
All people's tags	207	30%
(n = 686)		

Table 2. User pivot browsing behavior during field trial.

In addition to pivot browsing, the general search capabilities within *dogear* were also used by 30% of users (averaging eight search queries per user). The more restricted search forms (e.g., “by tag” and “by person”) were a late addition to the *dogear* service and so have we have observed negligible use to date.

Dogear User Survey

In addition to our analysis of actual *dogear* use, we have formally surveyed a sample of users of the *dogear* service. An invitation to participate in an electronic survey was emailed to 233 *dogear* users and 100 (43 % response rate). In general, the survey respondents were frequent *dogear* users; 44% used the service at least once a day and an additional 42% used the service weekly. 65% of the respondents self-reported that they were knowledgeable or expert users of other social bookmarking services. 43% of the respondents described their job function as IT specialist/IT architect, 27 % R&D, and the remainder spread over a variety of other IT-related jobs.

	Strongly Agree	Agree
Improved awareness of corporate resources or services	29%	49%
Improving my ability to find information on the corporate intranet	27%	45%
Increased awareness of the interests and expertise of other corporate employees	37%	34%
Better organization of my personal bookmarks	27%	39%
Increased ability to search and find web information	24%	42%
Improved my ability to find information on a specific topic	27%	37%
Decreased the time it takes to find a web resource	19%	40%
Increased sharing of information with my group or project team	17%	25%
Increased ability to search and locate someone with a specific interest or expertise	15%	26%
Decreased the time to find someone with a similar interest or skill.	4%	23%

Table 3. Survey results of perceived benefits of dogear

Of particular interest to us were the benefits provided by this social bookmarking service in performing various information seeking tasks. The percentage of respondents

who agreed with various task benefit statements can be seen in Table 3.

These survey results are generally quite favorable. A majority of respondents agree that *dogear* helps them to find information on both the corporate intranet and the external web. There are opportunities for service improvements, however, to increase the ability to share information with groups, and locate individuals with specific interest/expertise.

Other Dogear feedback

The *dogear* team has received feedback about the *dogear* service from two primary channels. There is an active and ongoing discussion of the *dogear* service on various intranet blogs. In addition, feedback is provided directly to the research team via email (contact information is provided). Much of the blog chatter is about announcements and evaluation of new features, and proposed feature enhancements. Occasionally usability feedback is offered.

Two of the more interesting topics are presented briefly here. The first is about issues around the idea of “social bookmarking” in general. One blog post said “*I'm still on the fence with all the social bookmarking thing... I mean, there is little time to go through what my todo list has listed, that I (up to now) find little value in either a) helping others get to what I find interesting, and b) using other people's tags to get to interesting stuff.*”...*On the flip-side, I do put a lot of value of the tagging approach to electronic documents (pictures, files, etc).*

Several responses to this post focused on the personal benefit of *dogear*. One user said “*Well, if you take out the social [shared] aspect of tagging, my biggest reason for switching is from an organizational standpoint. First, I find myself bookmarking a lot more sites that interest me, I may or may not return to them, but at least I put a little 'tag' on them and could always find them again relatively easily. Second, because I'm bookmarking so many more sites, hierarchical folders would be hard to manage, tagging makes it very easy to search/sort through my personal list. Lastly, the ability to access my bookmarks from any computer and any browser. So I see a number of personal benefits for using bookmark tagging, you could ignore the social aspect of them and still yield a lot of benefit to yourself*” A second commenter said. “*Indeed, I agree 100% with [name] on this one. Plus you also get the opportunity to subscribe to other people's (hopefully, with your same interests) bookmarks through RSS Feeds and also to the entire web site, whatever that may be. This is something very similar to what Furl and Spurl also do and offer.*” There are a number of other very supportive arguments for the benefits of a social bookmarking service within the company.

The second interesting blog “thread” is about standardization of tags for various reasons. One suggestion was to have *dogear* users tag bookmarks for new technologies with something like “learning.” The corporate learning group

could then subscribe to the “learning” RSS feed and present a summary of new bookmarks on the learning portal. Others have suggested getting corporate communities of practice (e.g., eclipse developers) to standardize on various tags so that information can be more easily shared among members. We plan to monitor these kinds of tag negotiations to understand how they influence the development of the *dogear* folksonomy.

Future Work: Exploiting the Links

One of the primary benefits of an enterprise social bookmarking application is that it is an explicit assessment of the utility or value of various internet and intranet information resources. As an information resource, there are many potential integration points with other corporate applications. Several of these are discussed here.

The first natural integration point would be to use social bookmark link structure to augment enterprise search applications. Web resources with high bookmark counts are likely to be relevant and useful. Indices of role or team-based bookmark collections can be specially weighted in team and role based search tools. An early variant of this search integration has already been developed as a Firefox plug-in. Intranet searches are augmented with a supplemental search of the bookmark collection, and bookmark matches are presented, as a group, in the search result list. If the extension recognizes the current browser location as a page of search results, it makes an HTTP request to the *dogear* server, and injects the resulting HTML into the page, in a manner consistent with the page design. We have created search-injection output formats for our own intranet search engine, and for Google. This appears to be promising as the intranet search injections have been observed 6003 times during field trial and 963 *dogear*-bookmarked links have been selected in preference to the results returned by the corporate search engine.

A second integration would be to combine an individual’s bookmarks with other information sources provided by that individual. A research prototype of an enhanced corporate directory service has already embedded portions of an individual’s bookmark collection in their directory information. The idea here is that an individual’s bookmark collection is meaningful reflection of the kinds of topics of general interest to them. This may be a powerful mechanism to provide dynamic updates to the “current interests” profiles for individuals. In a similar manner, corporate bloggers have embedded a list of recent personal bookmarks into their home blog page by using the JavaScript format provided by *dogear*.

A third natural integration point would be to combine bookmark collections from a group of individuals into web sites maintained to support various teams or projects. We have already observed one team, which is focused on e-learning, embed a collection of bookmarks on a team web site. Again, using the subscription services supported by the bookmarking application allows automatic updating of

bookmarks, promoting timely sharing of important information among team members.

CONCLUSION

While we have just begun to get feedback on the *dogear* social bookmarking service, we think that it shows great potential. There is a significant group of very active early users, who have already begun to generate excitement about the use of the application through informal communication mechanisms. Indeed, there are already over 100 mentions of the application in the enterprise blogs.

We have also already seen a handful of *dogear* extensions by other members of the organization, which will only serve to increase the benefits of the application. We are excited by the potential for using the *dogear* application to improve information sharing, expertise location and support of communities of interest within the enterprise.

We are encouraged by evidence that the *dogear* service provides bookmark management that meets both personal and organizational needs. The pivot browsing results show that individuals often explore their own bookmark collections, and both the survey results and blog posts indicate that *dogear* provides a means of improved personal bookmark management. The preponderance of shared (versus private) bookmarks, combined with the survey results, indicates that there is a willingness to share informational resources for the benefit of the organization. The substantial number of intranet bookmarks provide evidence that *dogear* is meeting a need for corporate social bookmarking.

At the time that this paper goes to press, the *dogear* service is still in use with a significant daily rate of browsing and new bookmarks. We will continue to monitor service use and solicit feedback from the user base. In addition to understanding the motivation for frequent use of *dogear*, we also will be investigating service abandonment rates (i.e. the number of people who try the *dogear* service but do not engage or continue to use the service).

Work is already underway to enhance the *dogear* social bookmarking service in several ways. As the number of shared bookmarks grow, so will the challenge to help users find new and particularly relevant content. We have begun to explore various search enhancements, as well as novel forms of customizable alerting. We have also begun to explore ways to provide better support for teams and communities of practice. These smaller groups may help us to understand what factors are important to “scale-down” the *dogear* service for small and medium businesses.

ACKNOWLEDGMENTS

REFERENCES

1. Abrams, D., Baecker, R., and Chignell, M. 1998. Information archiving with bookmarks: personal Web space construction and organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing*

- Systems* (Los Angeles, California, United States, April 18 - 23, 1998).
2. Buchanan, G. and Hinze, A. 2005. A generic alerting service for digital libraries. In *Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries* (Denver, CO, USA, June 07 - 11, 2005). JCDL '05. ACM Press, New York, NY, 131-140.
 3. Del.icio.us; <http://del.icio.us/>
 4. Fielding, Roy Thomas. Architectural Styles and the Design of Network-based Software Architectures. Doctoral dissertation, University of California, Irvine, 2000.
 5. flickr : <http://www.flickr.com/>
 6. Golder, S. and Huberman, B. A. 2005 The Structure of Collaborative Tagging Systems. Technical report, Information Dynamics Lab, HP Labs.
 7. Kanawati, R. and Malek, M. 2002. A multi-agent system for collaborative bookmarking. In *Proceedings of the First international Joint Conference on Autonomous Agents and Multiagent Systems: Part 3* (Bologna, Italy, July 15 - 19, 2002). AAMAS '02. ACM Press, New York, NY, 1137-1138..
 8. Keller, R. M., Wolfe, S. R., Chen, J. R., Rabinowitz, J. L., and Mathe, N. 1997. A bookmarking service for organizing and sharing URLs. In *Selected Papers From the Sixth international Conference on World Wide Web* (Santa Clara, California, United States). 1103-1114.
 9. Marais, H. and Bharat, K. 1997. Supporting cooperative and personal surfing with a desktop assistant. In *Proceedings of the 10th Annual ACM Symposium on User interface Software and Technology* (Banff, Alberta, Canada, October 14 - 17, 1997). UIST '97. ACM Press, New York, NY, 129-138.
 10. Millen, D. R. and Patterson, J. F. 2002. Stimulating social engagement in a community network. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work* (New Orleans, Louisiana, USA, November 16 - 20,
 11. My Web 2.0 Beta; <http://myweb2.search.yahoo.com/>
 12. Nonnecke, B. and Preece, J. 2000. Lurker demographics: counting the silent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands, April 01 - 06, 2000). CHI '00. ACM Press, New York, NY, 73-80.
 13. Technorati: <http://www.technorati.com/>
 14. Wasserman, S., & Faust, K..1994. *Social Network Analysis: Methods and Applications*. New York and Cambridge, ENG: Cambridge University Press.
 15. Weinberger, D. Taxonomies to Tags: From Trees to Piles of Leaves. *Release 1.0* February, 2005, pp. 2-32.
 16. Wittenburg, K., Das, D., Hill, W., Stead, L. 1995. Group asynchronous browsing on the World Wide Web, in *Proceedings of the 4th International WWW Conference*, (Boston MA, December 1995), O'Reilly & Associates, 51-62.