

SparTag.us: A Low Cost Tagging System for Foraging of Web Content

Lichan Hong, Ed H. Chi, Raluca Budiu, Peter Pirolli, and Les Nelson

Palo Alto Research Center (PARC)
3333 Coyote Hill Road, Palo Alto, CA 94304, USA

{hong, echi, budiu, pirolli, lnelson}@parc.com

ABSTRACT

Tagging systems such as del.icio.us and Diigo have become important ways for users to organize information gathered from the Web. However, despite their popularity among early adopters, tagging still incurs a relatively high interaction cost for the general users. We introduce a new tagging system called SparTag.us, which uses an intuitive Click2Tag technique to provide in situ, low cost tagging of web content. SparTag.us also lets users highlight text snippets and automatically collects tagged or highlighted paragraphs into a system-created notebook, which can be later browsed and searched. We report several user studies aimed at evaluating Click2Tag and SparTag.us.

Categories and Subject Descriptors

H5.2 [Information Interfaces and Presentation]: User Interfaces – Graphical User Interfaces. H5.3 [Group and Organization Interfaces]: Collaborative Computing.

General Terms

Design, Human Factors.

Keywords

Tagging, highlighting, social bookmarking, Web 2.0, annotation.

1. INTRODUCTION

Vannevar Bush's vision of the Memex [1] has inspired the evolution of information systems that augment and enhance human abilities to find, store, organize, understand, retrieve, and share knowledge. The areas of information retrieval, personal information management, and the Web (to name just a few) have, for the most part, historically been focused on supporting *individual* information foraging and sensemaking. Recently there has been an efflorescence of systems aimed at supporting *social* information foraging and sensemaking. These include social tagging/bookmarking systems for photos (e.g., flickr.com), videos (e.g., youtube.com), or web pages (e.g., del.icio.us). Tagging systems provide a means for users to generate labeled links (*tags*)

to content that, at a later time, can be browsed and searched. A unique aspect of tagging systems is the freedom that users have in choosing the vocabulary used to tag objects: any free-form keyword is allowed as a tag. Tags can be organized to provide meaningful navigation structures, and, consequently, can be viewed as an external representation of what the users learned from a page and of how they chose to organize that knowledge.

Empirical research [12] shows that people use tagging systems primarily for private, individual information storing and management. Other researchers [2,11,21] suggest that the success of social and collaborative systems is dependent on the architecture of interaction, as well as on the costs and benefits of interaction to the individual user. Grudin [11] discusses how costly interactions for the individual affect groupware adoption. Further, social tagging systems are instances of *networked information economies* [2], involving the production, distribution, and consumption of information by decentralized users operating over a network. Generally, as the costs of interaction are driven down, more users participate in the production and use of knowledge (e.g., tags). As more users participate, the value of the social tagging system to the participants is improved. Based on this reasoning, our research focuses on reducing the interaction costs of tagging.

In this paper, we describe a web content tagging system called SparTag.us. Our main objective is to provide low cost tagging and highlighting capabilities to users while they are reading web content. This paper has three main contributions:

- (1) Presenting a web content tagging system called SparTag.us that supports in situ tagging and highlighting during reading;
- (2) Introducing an intuitive *Click2Tag* technique to lower the overall costs of tagging; and
- (3) Showing (through user studies) that Click2Tag provides lower tagging costs than typing.

In Section 2, we present the need for low cost tagging and introduce the Click2Tag technique that allows users to tag a web page by clicking on words of that page. To validate the intuition behind Click2Tag, in Section 3 we analyze the provenance of tags in del.icio.us. After presenting the design and implementation of SparTag.us in Section 4, we report several user studies in Section 5. These studies investigated the impact of Click2Tag and collected user feedback on the usability of SparTag.us, as well as on its potential uses at the *individual* level. The studies were part of an iterative design process intended to improve the SparTag.us interface. Finally, we finish with some concluding remarks and future work.

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.

AVI'08, May 28–30, 2008, Napoli, Italy.

Copyright 2008 ACM 1-58113-000-0/00/0004...\$5.00.

2. MOTIVATION

The social bookmarking space has become quite crowded in recent years. Popular bookmarking systems include del.icio.us, MyWeb (myweb.yahoo.com), Diigo (diigo.com), Clipmarks (clipmarks.com), Magnolia (ma.gnolia.com), Bluedot (bluedot.us), and Google Notebook (google.com/notebook). Most of these systems allow tagging at either the page level (i.e., URL) or the sub-page level (e.g., individual paragraphs). A few of them, including Diigo, Clipmarks, and Google Notebook, also support other forms of annotation (e.g., highlighting).

To understand the costs of tagging, for each of these systems, we performed a GOMS-like analysis [6] of the interface and identified the overall number of steps involved in tagging. We count these steps to get a gross measure of the tagging costs, as shown in Table 1. Due to space constraints, we cannot discuss most of this analysis in detail. So, in the following, we exemplify this cost analysis for Diigo, since it offers features closest to what we had in mind.

Table 1. Tagging costs of social bookmarking systems.

System	Cost
del.icio.us	6
MyWeb	7
Diigo	8
Clipmarks	10
Magnolia	6
Bluedot	6
Google Notebook	11

To annotate a paragraph with Diigo, a user has to go through a sequence of steps as shown in Figure 1a. We identify two types of cost involved in this process: (1) *interaction cost* (mouse clicks, button presses, typing) in steps 1, 2, 3, 6, and 7; and (2) *attention-switching cost* (moving attention from one window to another) in steps 4 and 8, and possibly 5 (if the user has to go back to the original page to decide what tags to use). These costs reflect the interruption of reading and the switch to a different interface.

In SparTag.us, we aim to reduce these two types of cost by integrating tagging into the flow of reading. Specifically, with the Click2Tag technique we make two types of sub-page objects live and clickable: paragraphs of the web page and words of the paragraphs. According to our analysis of del.icio.us data (see Section 3), a considerable portion of tags comes directly from the web pages. When people tag a web page in del.icio.us, there is about 50% chance that the tag word appeared in the content of the page. Making each word of the paragraphs clickable allows users to simply click on a word to add it to the tag list without typing, thus lowering the interaction costs of tag typing. This input method can be especially useful in cases where keyboard input is not the primary input means (e.g., iPhones, tablet readers).

We also mitigate the attention-switching cost by enabling in situ tagging. With the Click2Tag interface, users can directly click on any paragraph to start tagging while they are reading it. When tags are entered, they are inserted at the end of the paragraph and displayed within the same rendered web page. This design was inspired by annotation studies [14,19] and by in-document collaboration systems [8], in which the content, and not some external form, is used as the setting for information sharing. Thus,

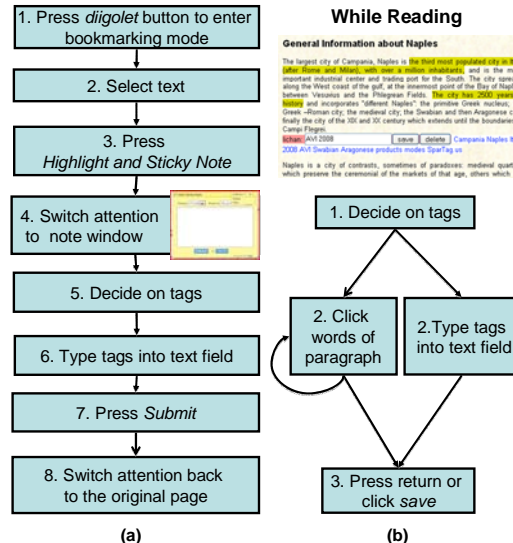


Figure 1. Cost analysis in (a) Diigo and (b) SparTag.us.

Click2Tag combines synergistically with in situ tagging, ensuring that the focus of attention remains on the content at hand.

Figure 1b shows a cost analysis of tagging that we strive to achieve in SparTag.us. Because users tag as they read, they do not need to shift attention to a dedicated tagging window, thus eliminating the attention-switching costs discussed earlier. We also eliminate the switching of attention back to the web page being read after the tags are applied. The attention-switching cost is reduced for those relying on consulting the document to decide which tags to apply (users get ideas for tags faster because the text is in front of their eyes). The interaction cost is also diminished: users can select tags from text instead of typing them. Moreover, compared to Diigo, we eliminated the initial button press since there is no need to invoke the tagging interface explicitly.

Studies have shown the need for richer forms of annotations due to the wide variety of annotation practices employed by readers [14,19]. Particularly for personal use and re-reading, Marshall and Brush [15] recognized the need for various forms of highlighting such as sentence highlighting, underlining, circles and margin bars. Many popular social tagging systems such as del.icio.us do not currently support other forms of annotations beyond tagging. In addition to tagging, SparTag.us also supports text highlighting. Indeed, our studies show that the combination of tagging and highlighting was perceived by users to be more valuable than either tagging or highlighting alone (see Section 5).

3. USING CONTENT WORDS AS TAGS

To further motivate the design of Click2Tag, we now present an analysis of how often tag keywords appear in the page content. Tagging is a process that associates keywords with specific content. This raises a question regarding the appropriateness of a content-driven approach such as Click2Tag: How many of the tags that people use actually come from tagged content? To help assess how Click2Tag relates to current practices employed in social bookmarking systems, we crawled del.icio.us and computed how often a keyword used by a user to tag an URL appears in the page content. We did this by the method described as follows.

First, we sampled tag activity data from the 29,978 most frequently bookmarked URLs up through June 2006. For each

URL, the body content was retrieved at the time of analysis (December 2007) if available and pruned to its textual content (i.e., the HTML tags were removed). Of the 29,979 URLs, we excluded 1,254 that had become invalid at the time of analysis. This left us with 28,724 valid URLs.

Next, for each URL, we converted its tag activity data into a set of tuples of the form (user, URL, tag). Each tuple corresponds to a single tagging event representing that a user bookmarked a specific URL using a particular tag. The contribution of a tuple is set to be 1 if the tag matches any word in the text body of the URL; otherwise, it is set to be 0. We computed the probability that a user tagged a page with a content word by first summing up the contributions of all the tuples for that specific URL and then dividing the sum by the number of tuples.

That is, for each URL, we collected all of the tuples involving that URL to give us a count of the total number of tuples for the URL. From this tuple collection, we then counted the number of tuples with content words as tags. Dividing the second count with the first count will give us the probability that a user tagged a page with a content word for that particular URL:

$$P(URL) = \frac{\# \text{ tuples - with - content - tags - for - URL}}{\# \text{ tuples - for - URL}}$$

By computing this probability for all of the 28,724 valid URLs, we get the distribution as shown in Figure 2 (see also the Color Page). Table 2 summarizes the data collection. On average, the chance that a tag comes from the content is 49%. This process produced a *conservative* estimate of tag occurrence in content,

Table 2. Summary statistics of del.icio.us data analysis.

# of URLs sampled	29,978
# of invalid URLs	1,254
# of valid URLs (searched for tags)	28,724
Probability that a user tagged an URL with a word occurring in page content:	
Mean	0.49
Standard deviation	0.26

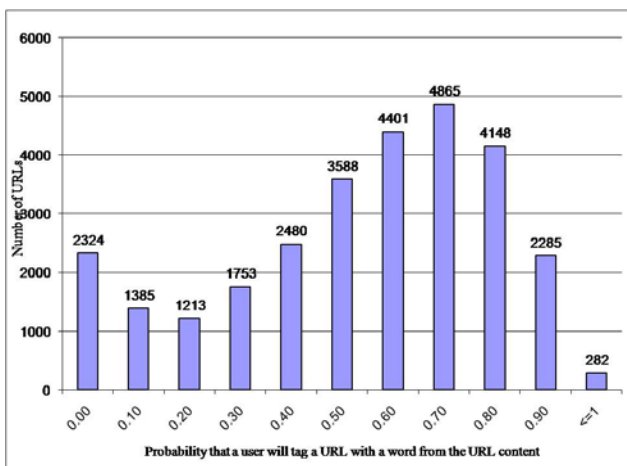


Figure 2. In del.icio.us, probability of a user tagged an URL with a word occurring in the page content of the URL.

since we did not account for situations such as content changes for a given URL (e.g., dynamic content), typos (e.g., “Ajaz” instead of “Ajax”), abbreviations (e.g., “ad” instead of “advertisement”), compound tags (e.g., “SearchEngine”), and tags written in languages other than that of the content.

To some extent, we were surprised by the relatively high probability of tag occurrence in del.icio.us tagged content. Clearly, people bookmark and tag pages for a wide variety of reasons [9]. While our analysis is not exact, it is reasonable to assume that when a word is used as a tag for a URL, there is a decent chance it derives from the content itself.

4. DESIGN OF SPARTAG.US

Having motivated the design of SparTag.us, we now describe its various functionalities and implementation in detail. In what follows, we use the word “annotation” to refer to both tagging and highlighting.

4.1 Lowering the Costs of Annotation

As mentioned earlier, in SparTag.us, we bring the tagging capability into the same browser window displaying the web page being read. When a user loads a web page in his browser, we augment the HTML page with AJAX code to make the paragraphs of the web pages as well as the words of the paragraphs live and clickable. Specifically, we identify the paragraphs by their HTML tags (e.g., elements enclosed within `<p>` or `<h2>`, etc). In addition, we alter the Document Object Model (DOM) tree of the page by enclosing each word of those paragraphs with the HTML tag ``. Furthermore, we attach mouse event listeners to the document object.

Our AJAX augmentation essentially converts each paragraph into a taggable object. As users read a paragraph, they can simply click anywhere on the paragraph to tag it. Consequently, a tagging widget is dynamically inserted at the end of the paragraph, as shown in the top part of Figure 3a (see also the Color Page). At this point, the user can input any tags into the text field of the widget. Clicking the *save* button or pressing the return key completes the tagging operation, with the tags displayed at the end of the paragraph (see Figure 3b). This Click2Tag interface provides a low cost way to invoke the tagging functionality while reading.

To add tags to the text field of the tagging widget, the user may simply click on words of the paragraph. Alternatively, the user can click on one of his most recently used tags (displayed in blue after the *delete* button in Figure 3a) to add it to the tag list. Of course, one can still type in the text field to add or modify tags, in case a desired tag does not appear in the paragraph. Recognizing the debate concerning single-word tags versus compound multi-word tags [13], we currently support single-word tags in SparTag.us, as in del.icio.us. To add compound tags in SparTag.us, the user needs to modify the existing tags in the text field or type directly.

In our augmentation, extra care is taken to handle special cases such as punctuations and hyperlinks. Punctuations are automatically filtered out. And by default, when a hyperlink is clicked, we let the browser follow the hyperlink to a new page. We use a word of the hyperlink as a tag only when the word is clicked with the control key held down.

The largest city of Campania, Naples is the third most populated city in Italy (after Rome and Milan), with over a million inhabitants, and is the most important industrial center and trading port for the South. The city spreads along the West coast of the gulf, at the innermost point of the Bay of Naples, between Vesuvius and the Phlegrean Fields. The city has 2500 years of history and incorporates "different Naples": the primitive Greek nucleus; the Greek –Roman city; the medieval city; the Swabian and then Aragonese city; finally the city of the XIX and XX century which extends until the boundaries of Campi Flegrei.

lichen: save delete Campania Naples Italy 2008 AVI Swabian Aragonese products modes SparTag.us

Naples is a city of contrasts, sometimes of paradoxes: medieval quarters which preserve the ceremonial of the markets of that age, others which are

(a)

The largest city of Campania, Naples is the third most populated city in Italy (after Rome and Milan), with over a million inhabitants, and is the most important industrial center and trading port for the South. The city spreads along the West coast of the gulf, at the innermost point of the Bay of Naples, between Vesuvius and the Phlegrean Fields. The city has 2500 years of history and incorporates "different Naples": the primitive Greek nucleus; the Greek –Roman city; the medieval city; the Swabian and then Aragonese city; finally the city of the XIX and XX century which extends until the boundaries of Campi Flegrei.

lichen: AVI 2008 Italy Naples Campania

Naples is a city of contrasts, sometimes of paradoxes: medieval quarters which preserve the ceremonial of the markets of that age, others which are

(b)

Figure 3. (a) Clicking on the first paragraph inserts a tagging widget to the end of the paragraph. (b) Tags are displayed as part of the first paragraph when tagging is finished.

The largest city of Campania, Naples is the third most populated city in Italy (after Rome and Milan), with over a million inhabitants, and is the most important industrial center and trading port for the South. The city spreads along the West coast of the gulf, at the innermost point of the Bay of Naples, between Vesuvius and the Phlegrean Fields. The city has 2500 years of history and incorporates "different Naples": the primitive Greek nucleus; the Greek –Roman city; the medieval city; the Swabian and then Aragonese city; finally the city of the XIX and XX century which extends until the boundaries of Campi Flegrei.

Naples is a city of contrasts, sometimes of paradoxes: medieval quarters which preserve the ceremonial of the markets of that age, others which are

Figure 4. Selecting a text snippet highlights it in yellow.

We can support highlighting fairly easily as well, since each word of the paragraphs is augmented with interaction handlers. To highlight a piece of text in yellow as shown in Figure 4 (see also the Color Page), the user first points the mouse at the beginning of the snippet. While pressing and holding the left button, he moves the mouse to the end of the desired text region and releases the button. This sequence of mouse movements is exactly the same as what it takes to select a piece of text in the browser, an operation that the user is already familiar with. Indeed, in our implementation we overload the text selection functionality of the browser to support text highlighting. We chose this implementation after several iterations of design, as discussed in Section 5.2. It also to some extent reflects the challenge of our AJAX augmentation.

At the end of the highlighting operation, we capture the highlighting region automatically and send the corresponding data to our web server asynchronously. In other words, the user does not have to press any *save* button to record the highlight. To allow the user to erase previous highlights and perform the default text selection operation, we create separate modes for highlighting, highlight erasing, and text selection, respectively.

Although in SparTag.us we mainly focus on paragraph annotation, it is conceivable that Click2Tag may be applied to

create tags at the page level as well. For example, a set of words selected by a sequence of clicks can be used as a basis to compute tags for the entire page. Compared to page tagging, paragraph annotation offers a better solution to handle pages that are fairly lengthy and yet only contain a few paragraphs of interest to the user. In addition, paragraph annotation serves as a visual indicator to remind the user what he feels is important on the page, which could be quite useful when the user revisits the page later on.

4.2 Notebook of Annotated Content

An important consideration in designing information foraging systems is to make it easy to store and retrieve the nuggets of information that have been collected by users [24,26]. Inspired by XLibris [23], we create a SparTag.us notebook for each user that automatically captures the results of the user's annotation activities. After the user annotates a paragraph on a web page, the paragraph is automatically extracted from the page and inserted into his notebook.

Figure 5 (see also the Color Page) shows the top portion of the notebook for user *lichen*. On the left side of the notebook, the annotated paragraphs are listed in reverse chronological order. For each paragraph, we show the time when the paragraph was last annotated. The tags and highlights created by the user appear in conjunction with the text of the paragraph. If needed, the user can directly modify his annotation on the paragraph here. At the end of the paragraph, we also include the URLs of pages that the user visited and contain the same paragraph. On the right side of the notebook, we show a tag cloud of the user's tagging keywords.

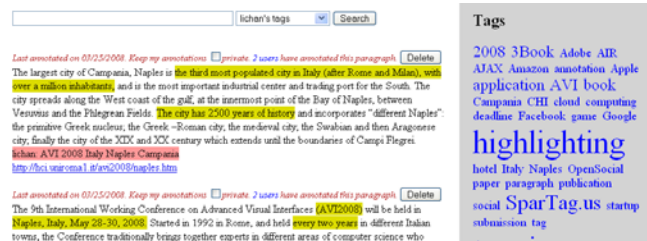


Figure 5. The top portion of the reading notebook that SparTag.us created for user *lichen*.

The search box at the upper-left corner of the notebook provides several search options. The user can search against the paragraph tags, the highlighted snippets, the entire paragraph texts, or the URLs. This makes it easy to retrieve a subset of paragraphs that are related to a certain concept or from a specific web site.

4.3 Architecture

In our implementation, SparTag.us consists of two parts: a client-side browser extension and a server. The client side is currently implemented as a Firefox extension, which includes a browser toolbar. The toolbar, as shown in Figure 6, consists of several shortcuts to key features of SparTag.us, e.g., turning on/off SparTag.us, changing modes, and accessing the notebook. We are currently exploring client extensions for other browsers as well, including IE and Safari.

The client side extension also contains a GreaseMonkey [10] script. GreaseMonkey is a client side tool that inserts our AJAX code into web pages while allowing the user's browser to communicate with multiple web servers. Figure 7 shows an architecture diagram of SparTag.us. As shown here, after a web

Figure 6. Browser toolbar of SparTag.us.

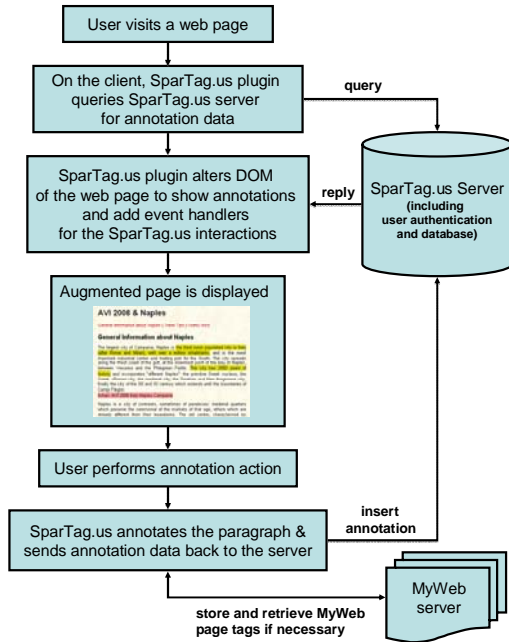


Figure 7. Architecture diagram of SparTag.us.

page is loaded onto Firefox, the GreaseMonkey script extracts the paragraph texts and forwards them to the SparTag.us server. In return, it retrieves annotation data created by the user for the paragraphs. Subsequently, GreaseMonkey alters the DOM of the page to display the annotations and control its event handling. When the user tags or highlights a paragraph, the script submits the annotation data to the SparTag.us server, which is then stored in a database.

The SparTag.us server is an Apache Tomcat server that runs Java servlets and connects to a MySQL database. The database stores the user’s paragraph tags and highlights. For example, a tag entry includes things such as what the tags are, which paragraph the tags are attached to, who created the tags and when. The server also offers other services such as user authentication. In the spirit of mashups, we also support page tagging, using Yahoo’s MyWeb to store and retrieve page tags. We created this mashup to offer both paragraph and page taggings to the user as well as to take advantage of MyWeb’s vast user community and its services.

5. EMPIRICAL STUDIES

In conducting our evaluation of SparTag.us and Click2Tag, we were interested in two issues. First, we wanted to understand whether Click2Tag was indeed a low cost technique, when compared to the traditional way of tagging by typing. Second, we needed to look at the system usability and at how the system is used by real users in everyday practice.

5.1 Evaluation of Click2Tag

We conducted a behavioral experiment to study how Click2Tag compared to type-to-tag or no-tagging. In this experiment, we were interested in the relative costs of Click2Tag and type-to-tag. Our study was part of a more complex experiment that examined

comprehension and memory effects. As we shall mention later, a preliminary eye-tracking study showed that processing is quite different in the two conditions, and it led us to also hypothesize that the two tagging techniques would have different effects on comprehension and memory. Due to space constraints, we only present the findings related to encoding costs and just briefly mention the other memory-related results.

Participants. We recruited 27 participants for this study. Each participant was compensated with \$20.

Materials. We selected 18 passages from news articles as well as from various web pages on the Internet. The passages reflected a variety of topics (medicine, education, general science, aviation, history, etc). On average, the passages were 267 words long (ranging from 253 to 279).

Procedure. The study was structured as a within-subject study. Participants had to perform 18 trials (each corresponding to one of the 18 passages). In each trial, participants read a passage, selected randomly from the list of 18 passages. Participants were instructed to read at their own paces, but if they spent more than 2 minutes on a trial, they were moved to the next trial. They were also told to try to retain as much content as they could from the text (and they had to remember that content in a different part of the experiment that we do not discuss here). The trial could belong to one of three interface conditions as follows:

- (1) **No-tags:** In this condition, no tagging was performed.
- (2) **Click2Tag:** Participants had to tag the passage with relevant words by clicking on words from the passage. The tags were displayed in a box under the passage and could not be modified by the participants.
- (3) **Type-to-tag:** Participants had to tag the passage with any relevant tags that they could generate, and type those tags in a box under the passage.

Results. We performed ANOVAs with subjects as the random factor using the interface condition (no-tags, Click2Tag, type-to-tag) as an independent variable. Table 3 shows the average reading time and number of inputted tags per condition. There was a significant effect of condition ($F(2,52)=52.72, p<0.001$). Contrasts showed that participants spent less time in the no-tags condition than in the Click2Tag condition ($p<0.001$) and in the type-to-tag condition ($p<0.001$), and also participants in the Click2Tag condition were faster than in the type-to-tag condition ($p=0.05$). These results pointed to a time cost associated with the tagging conditions, and confirmed that Click2Tag was a lower-cost interaction than type-to-tag.

People tended to attach more tags ($p<0.001$) in the Click2Tag condition (6.95 on average) than in the type-to-tag condition (3.98 on average), suggesting that they took advantage of the ease to tag in the Click2Tag condition to attach more tags and they achieved this faster than in the type-to-tag condition. These results indicate that, since Click2Tag is faster per tag, people may end up with more tags per passage than in the type-to-tag condition. The implications of this effect on information retrieval and formation of tag folksonomies still remains to be explored.

Table 3. Reading times and number of tags.

	Reading Time (seconds)	Number of Tags
No-tags	81.38	0
Click2Tag	96.01	6.95
Type-to-tag	102.92	3.98

Implications for memory and learning. In the following we briefly review the implications for memory and learning for the two tagging mechanisms. Figure 8 (see also the Color Page) shows hotspot gaze maps for a participant going through a type-to-tag trial versus a Click2Tag trial. As suggested in the Click2Tag condition (see Figure 8b), people fixated more on particular words in the text. This appears to have led them to rehearse the text content more and to perform bottom-up, content-driven tagging. In contrast, in the type-to-tag condition (depicted in Figure 8a), people tended to fixate less on any specific words, possibly because they read the text and then used their background knowledge to generate tags for the text, as opposed to fixating on specific words. Thus, the tagging process was top-down, knowledge driven and entailed more elaboration and connection with prior knowledge.

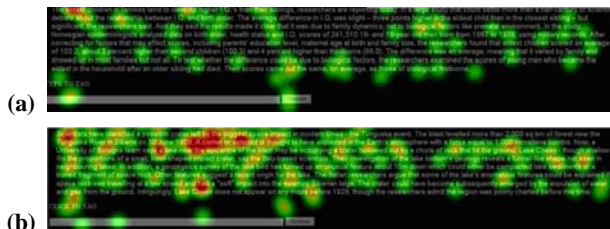


Figure 8. Hotspot gaze maps for two tagging techniques: (a) type-to-tag and (b) Click2Tag.

As a consequence of these two different processing styles, one would expect to obtain better retention of individual facts from text in the Click2Tag condition (since practice is related to recognition memory [22]), but better free recall in the type-to-tag condition (since elaboration improves recall [3]). Our data confirms that Click2Tag led to better recognition memory (both in terms of accuracy and time to recognize the facts) than type-to-tag. However, although there was a tendency for the type-to-tag condition to facilitate recall, this was not significant. In [5] we discuss these effects in more detail and argue that this lack of effect of type-to-tag on recall is due to the high cost of tagging by typing. Because they have to spend time typing rather than studying the text, people lose whatever memory gains they may get by having engaged in elaboration to find tags.

The memory data delivers one more piece of good news for Click2Tag: the recall efficiency of Click2Tag (i.e., number of facts recalled per unit of study time) was significantly better than that of type-to-tag, but not different from the recall efficiency of the no-tags condition. This result suggests that for type-to-tag, the supplemental cost of typing affects recall. While in Click2Tag, due to its low cost, the effects on recall are comparable to the no-tags condition, since most of the time is spent on reading.

5.2 Usability and User Feedback

We also conducted two separate user studies as part of an iterative design process to address how usable SparTag.us was and how it could be improved. The first study was conducted in the lab and consisted of people using SparTag.us to perform specific tasks, and then giving us feedback. In the second study, we had several employees of our company use SparTag.us for an extended period of time and we conducted in-depth interviews with some of these users to collect general feedback.

5.2.1 Usability Study

Participants. The users in our study were 10 company employees, older than 25, who volunteered their participation.

Procedure. Users were asked to complete six tasks. In the first three tasks, they used SparTag.us to find and tag information from the Web (e.g., *find a good pair of headphones under \$200*). The other three tasks required the users to use the SparTag.us notebook to retrieve information that they may have collected or encountered during the first three tasks (e.g., *find good ear bud headphones under \$100*). We videotaped all the tasks. At the end of the study, users filled in a 41-item survey including the SUS usability questionnaire [4] and other questions derived from Nielsen’s 10 usability heuristics [17], intended to help us understand the usability of SparTag.us. We also asked questions related to system features, and invited open-ended comments.

Results. According to the post-study survey, participants found SparTag.us intuitive and easy to use and felt that they could use the system without having to read the help page. With regard to usability, SparTag.us was rated as a moderately usable system. The average SUS usability score was 64 (ranging from 45 to 82.5, standard deviation 13.13). For comparison, Windows ME has a score of 55.73 [16]. Average scores for SUS are typically between 65 and 70. The following are several main lessons that we learned from the users’ feedback:

- Most users considered Click2Tag and the low cost interface as the main benefit of SparTag.us (e.g., *[I liked the] easy addition of tags without extra pop-ups and the easy display of tags overlapped over a page*).
- Paragraph tagging (e.g., *[I liked] annotating in-depth, not stuck with just a label for the entire content, which is a problem of most tagging systems*) as well as easy search and retrieval using the notebook were other mentioned benefits.
- Users appreciated SparTag.us for personal data organization (*[I liked the] ability to capture and organize information. It was easy to learn and intuitive, and it appealed to my personal interest*). Indeed, 7 out of 10 users said they would use SparTag.us primarily for the purpose of organizing personal information. Moreover, 9 out of 10 participants said that, when they revisited a page, it was useful to see the annotations they had made on previous visits.
- Due to our initial design, users felt that tagging was easier than highlighting and that highlighting was too hard to perform without making mistakes.
- Even though the highlighting design had problems, most users (70%) felt that the combination of tagging and highlighting was most useful, rather than only tagging or only highlighting.

Improvements made. A close inspection of the videos taken in the study revealed that highlighting provided little feedback to the participants on whether the highlighting interactions were successful. In the next version of SparTag.us, we created separate modes for highlighting, highlight erasing, and text selection, respectively. We also observed that the participants experienced difficulties in making precise mouse clicks to define the highlighting boundaries. We resolved this shortcoming by tracking where the mouse enters and exits a paragraph, and updating the highlight as the user moves the mouse. This mirrors the familiar text selection interactions.

5.2.2 Assessment by Field Trial

Our experience in the laboratory suggested that SparTag.us streamlines a reader's entry of annotations to web content. We wished to examine how this works in people's everyday practices. Consequently, we conducted a 6-week field trial of SparTag.us with 18 people in our own organization outside of the design team. This allowed close support and interaction with users in one social setting, but also exposed the new technology to some range of diversity. Participants ranged in role from summer interns to senior staff and managers, working within four different areas of our organization. The development team stated no constraints or guidance other than feature description and support.

Our initial assessment of use practices is through in-depth interviews with selected users. Based on the observed frequency of use, we chose four participants for discussions conducted at least several weeks after initial deployment. We selected four people who gave the system an extended trial (above average in terms of pages viewed with SparTag.us running). Three of the participants (Interviewees 1, 2 and 3) were above average users of the technology in terms of the amount of annotations made, and one (Interviewee 4) was below average.

Interviews were conducted by a researcher without prior involvement in the project, and were semi-structured, loosely following 26 pre-planned questions, but allowing the interviewee to cover the material as the topics arose in the interview. We designed the questions to walk the person through each of the various components and features of the SparTag.us interface, and phrased them to elicit stories about instances of application use (e.g., *"Could you tell me about the last time you used SparTag.us? When was that? What were you doing? What happened?"*). These questions were then followed up for as many different types of instances as the person could recall. Each interview was conducted in the office of the user and audio-video recorded. Participants were allowed to use the interface to demonstrate actions and look up information in the system to refresh their memories of events.

Reported benefits. People did engage with tagging at the paragraph level (in situ) and utilized the highlighting feature to mark salient points. As Interviewee 3 said while talking about a computer operation web page, *"That part [of the document] was interesting to me, not so much the other parts. But this is something I spend some time looking for and I really want to get a handle. I highlighted that. Oh, yeah, then I put some tags in, dual boot, sort of reminding myself. Remind myself that this is about partitioning and dual booting."*

The focused annotation on specific subsets of a document, and the ease of tagging and capturing paragraph text in the notebook without extra effort was mentioned as what interviewees like best about the system. Interviewee 2 said, *"[...]with ClipClip I need to type, notes, to remind myself, type the tags. I don't like that";* and *"So I like the notebook page, it gives me at a glance, like the specific content I was interested in as well as the URL. So it serves the same purpose that I use del.icio.us for because it saves the URL, but it gives me more at a glance the interesting content."*

Reported limitations. All users interviewed reported turning the system on and off (and then forgetting to turn it back on for a while). Much had to do with usability and performance of an early prototype. Of greater interest was that all users reported accidentally tagging things. Each developed the workaround

strategy of abandoning the page (and hence not saving the erroneous tags). This points out that, while the Click2Tag implementation was described as easy to invoke, not all people will wish to invoke it in the same way with respect to their own browsing practices (Interviewee 4, *"I had to tell it to stop highlighting things. I would prefer it to be much more in the background and become available when I want it"*).

Improvements made. From the findings of our field trial, we improved the performance by reworking the SparTag.us client implementation. We also discovered that our highlighting implementation had led to problems with some pages containing HTML forms. To address this issue, we overloaded the text selection functionality of the browser, using the browser's selection object to define the highlighting region. To enable users to find out the status of SparTag.us, turn it on if not, and repair tagging or highlighting miscues, we developed the SparTag.us toolbar, as described in Section 4.3. With one mouse click, users can now turn on or off SparTag.us and see the effect immediately. The toolbar also shows which mode (highlighting, highlight erasing, or text selection) SparTag.us is currently in, and the mode can be changed interactively. Furthermore, we added additional feedback and correction mechanisms. For example, users can easily see what text is annotatable and undo a tagging or highlighting operation with a mouse click or a shortcut key.

6. CONCLUSION AND FUTURE WORK

We have described a novel tagging system called SparTag.us that introduces a new technique, Click2Tag, to support low cost, in situ tagging. One motivation for the Click2Tag technique came from our analysis of del.icio.us data, which showed that a large number of tagging events involved applying tags that came from the page content. Another motivation was that tagging in current social bookmarking systems has relatively high interaction and attention-switching costs. Our first user study examined how the cost of Click2Tag compared with that of typing to tag. The results suggested that people are more efficient with Click2Tag. They process the text faster and tend to attach more tags.

SparTag.us also allows in situ highlighting and automatically stores all the paragraphs that a user has highlighted or tagged into a notebook. We described two user studies addressing the usability of SparTag.us. These studies have been informative in making interface design decisions. They also reassured us of the potential user acceptance of our ideas on Click2Tag, the combination of tagging and highlighting, and the notebook. Obviously, Click2Tag is designed for web pages whose contents are mainly textual, not for pages consisting of media contents (e.g., images and videos).

There are many issues to be addressed in our future work. First, we need to compare the quality of tags generated by Click2Tag versus typing, as well as their effectiveness in information search and retrieval. In a recent paper [5], we reported the different implications of these two tagging techniques on memory and learning, which shed some light on the potential effects on later retrieval using these tags. Moreover, the folksonomies arising from these two techniques are likely to be different, since, as we discuss in Section 5.1, Click2Tag is content-driven, whereas typing is more knowledge driven.

Second, another question of interest is how in situ tagging and highlighting affect within-page navigation. If you come back to a

page and see your own (or others') annotations, will you be able to retrieve content of interest faster? Generally, we believe that tagging and highlighting support different functions, although there are certainly some overlaps. Tagging seems to help with navigation mostly, whereas highlighting seems to support learning either at encoding time or at reviewing time [7,18,20,25]. Future research is needed to understand whether the combination of tagging and highlighting will offer more support for gist reconstruction than tagging alone.

Third, we plan to develop the social aspect of SparTag.us. In this paper we focused the description on the individual usage of SparTag.us, under the assumption that we first need to make the individual experience as effective as possible in order to be able to attract users to the system. Indeed, despite various burrs to be expected of the early prototypes, users found our system to be as usable as existing well-developed systems. The next step of our research is to explore the design space of social functionalities. We are interested in finding out what kind of features would support social interaction and communication (e.g., seeing your friends' tags or notebooks, being able to share tagged or highlighted content via email or other means, etc).

Lastly, we are interested in exploring the Click2Tag concept on documents other than web pages (e.g., PDF and MS Word documents). Since for most people web browsing is hardly the only means to acquire new knowledge, we plan to investigate how to share annotations across documents of different formats.

7. ACKNOWLEDGMENTS

We thank the user study participants, Lawrence Lee for helpful discussion, and Christiaan Royer for help in experimental setup.

8. REFERENCES

- [1] Bush, V. (1945). As We May Think. *The Atlantic Monthly*, 176(1), 101-108.
- [2] Benkler, Y. (2002). Coase's Penguin, or Linux and the Nature of the Firm. *Yale Law Journal*, 112, 367-445.
- [3] Bradshaw, G. and Anderson, J. (1982). Elaborative Encoding as an Explanation of Levels of Processing. *Journal of Verbal Learning and Verbal Behavior*, 21, 165-174.
- [4] Brooke, J. (1996). SUS: A "Quick and Dirty" Usability Scale. In P. W. Jordan, B. Thomas, B. A. Weerdmeester, and A. L. McClelland (eds.), *Usability Evaluation in Industry*. Taylor and Francis, London.
- [5] Budiu, R., Pirolli, P., and Hong, L. (2007). Remembrance of Things Tagged: How Tagging Affects Human Information Processing. Under submission.
- [6] Card, S., Moran, T., and Newell, A. (1983). *The Psychology of Human Computer Interaction*. Lawrence Erlbaum.
- [7] Chi, E., Gumbrecht, M., and Hong, L. (2007). Visual Foraging of Highlighted Text: An Eye-Tracking Study. *Proc. HCI International Conference*, 589-598.
- [8] Churchill, E., Trevor, J., Bly, S., Nelson, L., and Cubranic, D. (2000). Anchored Conversations: Chatting in the Context of a Document. *Proc. CHI'00*, 454-461.
- [9] Golder, S. and Huberman, B. A. (2006). Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science*, 32(2), 198-208.
- [10] GreaseMonkey. DOI=<https://addons.mozilla.org/en-US/firefox/addon/748>.
- [11] Grundin, J. (1994). Groupware and Social Dynamics: Eight Challenges for Developers. *Communications of the ACM*, 37(1), 92-105.
- [12] Lee, K. J. (2006). What Goes Around Comes Around: An Analysis of del.icio.us as Social Space. *Proc. CSCW'06*, 191-194.
- [13] Linderman, M. (2005). Tag Formats: Can't We All Just Get Along? DOI=http://37signals.com/svn/archives2/tag_formats_cant_we_all_just_get_along.php.
- [14] Marshall, C. (1997). Annotation: From Paper Books to the Digital Library. *Proc. Digital Libraries'97*, 131-140.
- [15] Marshall, C. and Brush, A. (2004). Exploring the Relationship between Personal and Public Annotations. *Proc. Digital Libraries'04*, 349-357.
- [16] Microsoft. DOI=http://download.microsoft.com/download/d/8/1/d810ce49-d481-4a55-ae63-3fe2800cbabd/ME_Public.doc.
- [17] Nielsen, J. (1994). Heuristic Evaluation. In J. Nielsen and R. L. Mack (eds.), *Usability Inspection Methods*. John Wiley and Sons, New York, NY.
- [18] Nist, S. L. and Hogrebe, M. C. (1987). The Role of Underlining and Annotating in Remembering Textual Information. *Reading Research and Instruction*, 27(1), 12-25.
- [19] O'Hara, K. and Sellen, A. (1997). A Comparison of Reading Paper and On-Line Documents. *Proc. CHI'97*, 335-342.
- [20] Peterson, S. E. (1992). The Cognitive Functions of Underlining as a Study Technique. *Reading Research and Instruction*, 31(2), 49-56.
- [21] Pirolli, P. (2007). *Information Foraging Theory: Adaptive Interaction with Information*. Oxford University Press, New York, NY.
- [22] Pirolli, P. and Anderson, J. (1985). The Role of Practice in Fact Retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 136-153.
- [23] Schilit, B., Golovchinsky, G., and Price, M. (1998). Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. *Proc. CHI'98*, 249-256.
- [24] Schraefel, M. C., Zhu, Y., Modjeska, D., Wigdor, D., and Zhao, S. (2002). Hunter Gatherer: Interaction Support for the Creation and Management of Within-Web-Page Collections. *Proc. WWW '02*, 172-181.
- [25] Silvers, V. L. and Kreiner, D. S. (1997). The Effects of Pre-Existing Inappropriate Highlighting on Reading Comprehension. *Reading Research and Instruction*, 36(3), 217-223.
- [26] Thomas, J. and Cook, K. (2005). *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE CS Press, Los Alamitos, CA.