Lightweight Tagging Expands Information and Activity Management Practices

Gerard Oleksik[†], Max L. Wilson[ϟ], Craig Tashman[‡], Eduarda Mendes Rodrigues[*], Gabriella Kazai[*], Gavin Smyth[*], Natasa Milic-Frayling[*], Rachel Jones[†]

[†]Instrata
12 Warkworth Street
CB1 1EG
{geoleksik, racheljones}
@instrata.co.uk

[ϟ]Future Interaction Technology Lab
University of Swansea
Swansea, SA2 8PP
M.L.Wilson@swansea.ac.uk

[‡]Georgia Institute of Technology
GVU Centre, Atlanta
GA 30332
craig@cc.gatech.edu

[*]Microsoft Research
7 J J Thomson Avenue
CB3 0FB
{eduarda, gabkaz, gavin.smyth, natasamf}
@microsoft.com

ABSTRACT

Could people use tagging to manage day-to-day work in their personal computing environment? Could tagging be sufficiently generic and lightweight to cover diverse working practices and even support new practices for managing applications and accessing documents efficiently? We investigate these issues by implementing the TAGtivity system that enables users to tag any (or many) type of resources in the context of their ongoing work. We deployed TAGtivity and studied users' tagging practices in actual work places over a three week period. Our analysis of interviews and logs reveals that affordances of the TAGtivity system supported users' information and activity management practices. These include new practices for managing emerging activities and ephemeral information and for optimizing access to documents across application data silos.

Author Keywords

Tagging, activity management, information management, user evaluation.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

In 1983, Bannon et al. [1] suggested that electronic resources used in day-to-day work should be managed in relation to the user's activities. Since then there have been many attempts to apply this principle to assist users in managing their documents and applications [2, 5, 17].

From the research literature, we draw a distinction between supporting users in managing their applications and application windows, often referred to as *activity management*, and organizing and accessing documents within the file system or specialized content management systems (e.g., email), i.e., *information management*. Activity management includes handling of multiple application windows, switching between tasks, managing interruptions, and preserving the context of their work. Information and file management typically refers to organizing resources, i.e., files, folders, emails, web pages, and the like, for easy access, publishing, and sharing.

Studies have shown a wide disconnect between the user's organization of the file system and the information access the user requires during everyday work practices [2, 5]. Users often need quick access to resources from potentially disparate parts of the file system. In some instances, relevant information is associated with proprietary information stores that cannot be accessed easily, except through the application or service itself. Examples include e-mail services, Web resources, and bookmarks managed within Web browsers.

Recent projects explore the use of semantic tags to label documents and thus provide alternative ways of organizing and accessing documents [4]. While community tagging services, such as *Flickr* and *Del.icio.us*, have been studied extensively, we lack in-depth analysis of resource tagging within the PC environment.

Our work helps bridge that gap. It includes: (1) design and implementation of a generic tagging system, TAGtivity, for tagging resources within a PC environment, (2) in-situ study of tagging practices, comprising the deployment of the TAGtivity system, logging of user's activities, and user interviews, and (3) in-depth analysis of the collected data. In preparation for the study we invested considerable effort in designing new and flexible tagging support but our primary objective was to observe and characterize the emerging tagging behaviours rather than evaluate the effectiveness of individual features or design options. Thus, our main contributions stem from the user study and the insights we gained from the data on how, when and why users create tags, and how that relates to their broader work practices.

CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.

Our research reveals that tagging extends the utility of the file system by providing additional views or *logical organization* of the content included in the static file organization. Furthermore, the tagging enables capture of ephemeral information that would not warrant inclusion into permanent folders of the file system. Finally, the TAGtivity system aids activity management in several ways; by using tags to collect resources related to a task, by enabling flexible switching between tasks, and by enabling association of resources to multiple tasks.

In the following sections we reflect on related work and provide a description of the TAGtivity system. We then discuss in detail the study design and methodology. In the core part of the paper we present in-depth analyses of the study data and the study findings. We conclude with the discussion and summary of our work.

RELATED WORK

Our literature review is focussed on research that deals with file management issues, tagging practices, and the design of systems for activity management.

File Organization

Information management in PC environments has long been dominated by the hierarchical folder metaphor. While this system offers many benefits to the user, such as bearing a resemblance to a real-world analogue, it also suffers from disadvantages, as highlighted by Hsieh et al. [9] and by Golder and Huberman [6]. These include a potentially high cognitive load for memorising hierarchies, particularly challenging for large number of folders that users frequently possess, and the tendency to forget information that is 'out of sight' [10].

Jones et al. [10], for example, studied the meaning and structure of the folder hierarchies amongst 14 users. They highlighted the number of recurring folder names and organisational structures that stemmed from the user practice to use the same folder structure from project to project. Whilst Jones et al. [10] argue that such behaviour could be supported through better tools for cloning existing folder structures, this also suggest an alternative organizational system that allows files to be organized along multiple, orthogonal dimensions at once. Folders cannot provide this form of organization, as they are based on a location metaphor. A document is found by returning to its *location* in the folder hierarchy. As such, it is inconsistent for the same document to be in multiple, non-nested folders at once. Organization based on tagging avoids this; documents may coherently possess any combination of tags. Our work builds on this premise with the aim to aid users in managing their resources across multiple activities.

Tagging Practices

Tagging has been applied as an organizational and classification scheme in a variety of systems. As noted previously, services such as *Flickr and Del.icio.us* allow users to tag either photographs or web links, aiding the retrieval and organization of these resources [6]. Through these and similar online services, tagging has become popular as a tool for content browsing and discovery. Recently, it has also made its way to the PC environments, e.g., through the tagging features of Microsoft Windows Vista, and complementary applications such as *VistaGlance*[1], which enable users to tag their documents.

A tagging tool for the PC is Phlat [4]. This system facilitates document retrieval by allowing the user to tag files, emails, calendar entries, and the like (but not Web pages). Phlat was deployed with a large number of users, reporting on the statistical analysis of its usage. However, no substantial qualitative findings have been made available to the scientific community. In contrast to [4], we designed the TAGtivity to include tagging of Web pages and focussed on the qualitative analysis of the tagging practices that emerged during our study.

Conceptually the closest to the TAGtivity approach are Giornata [17], the Placeless Documents project, and the closely related Presto system, by Dourish et al. [5]. These systems include activity and content management based on tagging. Presto, for instance, allowed users to apply user specified attributes to documents and use them to retrieve, index, and organise documents into 'fluid collections' that support specific tasks. Interaction with these collections was facilitated through *Vista*, a browser which allowed users to view collections and further add attributes to documents. However, the tagging facility was not closely integrated into the UIs of the desktop applications, as we achieved in the design of the TAGtivity system. Furthermore, Presto was not evaluated through a user study and thus leaves open questions about how users would adopt tagging to manage their activities.

Hsieh et al. [9] present a web-based tool for organizing personal documents. They draw upon cognitive psychology to argue that tagging in the personal information space offers a better fit with the workings of the human memory than hierarchically organized folders do, further suggesting that tagging may be a valuable addition to traditional hierarchical organization methods.

Although tagging has been applied through a variety of systems, there have been few studies on the use of tags in the realm of personal document or resource management. There is also a lack of empirical data offering insights into the nature of tag creation and use, and the motivation for using tags. In our study we aim to address this gap and provide both qualitative and quantitative analysis of the study data collected through interviews and logs recorded by the deployed TAGtivity system.

Activity Management

A number of systems and approaches have been developed to manage applications and documents that are actively

[1] http://vistaglance.com – Home – Vista Glance

used for user tasks. While the specific qualities of these systems vary significantly, they share a common objective to help the user group and manage related application windows. They differ primarily in the way they represent the groups of windows and the affordances by which the users can create and manage them.

Virtual Desktop Management (VDM) constitutes one approach to activity management. Drawing on the concept first introduced by Henderson and Card's Rooms system [8], VDMs divide the user's environment into a number of virtual desktops (rooms) that can be used to separate the resources that are associated with distinct activities. The user can switch between activities by moving between rooms. While Rooms allowed resources to exist in multiple desktops, other manifestations of VDM, such as Task Gallery [13] and Kimura [11], do not.

Giornata [17] also takes a VDM approach to activity management but incorporates tagging as well. Users can tag individual desktops and any file accessed within a particular desktop is automatically linked to the corresponding desktop tags. In addition, Giornata enables users to assign individual tags at the file level through the MacOS file properties window. However, such an action would not be conducive to lightweight tagging since the user would have to manually open and edit properties each time they wanted to create or modify tags. Finally, Giornata does not use tags as a means of retrieving files and windows, which are still organized in a traditional VDM manner.

Thus, we concur that Giornata and TAGtivity show some similarity in features. However, they have been designed for different purposes. Giornata is focussed on activity management while TAGtivity is designed for generic and flexible tagging of resources.

Giornata was deployed with 5 participants, who used it in their everyday work for an average length of 54 days. While the authors report that participants' reactions to the system were positive, they do not present an in-depth discussion of participants' use of the system.

Two other notable approaches to activity management are represented by GroupBar [14] and Scalable Fabric [12]. In GroupBar, proxies (taskbar buttons) representing application windows can be dragged together to form a group. The user may then show or hide entire groups of windows to facilitate switching between activities. Scalable Fabric allows windows to be miniaturized and then grouped together on the desktop. The user may selectively expand or miniaturize these groups as they change activities. However, while both of these systems allow lightweight window groupings, they do not allow windows to be associated with multiple groups at once.

Another approach, conceptually similar to VDMs, is explored by Bardram et al. [2], whose Activity Based Computing (ABC) framework enables activity management as well as roaming and collaboration across the PC environments. In ABC activities are created through a centralised activity bar which allows users to aggregate resources into groupings that could be resumed or suspended in order to switch tasks. While system evaluation revealed it to be useful and easy to use, it also revealed several problems. The first is related to the lack of support for simultaneous use of the same resource in multiple activities. The second refers to the mismatch between the system design and the activity life cycle, in particular with respect to emerging tasks. The issue is contingent emergence of activities where multiple activities may begin to overlap in complex ways. For example, whilst in one activity the user may open a new window, which may potentially pertain to a new activity. Bardram's solution was to allow disassociated windows from the current task to remain open during the task suspension, allowing these windows to form the basis of a new activity.

Gonzales et al. [7] have also drawn attention to the emergent nature of tasks and, through diary studies, shown that unexpected interruptions were a common source of new activities in office work. This suggests that activity management needs to support disruptions. Unfortunately the above systems generally lack the flexibility to support emergent activities.

As with the hierarchical folders, much of this difficulty can be traced to the use of location-based metaphors for representing groups of resources. One exception is WindowScape [15]. Like Scalable Fabric, WindowScape enables windows to be represented as shrunken miniatures but uses a temporal rather than spatial metaphor for representing groups. This approach does enable windows to be associated with multiple groups simultaneously but faces a scalability issue as the user's interaction history grows. The system is also focused on managing windows rather than general resource tagging as in case of TAGtivity. An alternative approach to activity management is exemplified by [3] whose email based Taskmaster system brought task management tools into the inbox.

In conclusion, while there have been previous attempts to use tags for organizing and accessing information and managing tasks, our work is among the first to deploy a flexible tagging approach that applies to both problem areas and enables us to derive insights from observed user practices. In the next section we describe the TAGtivity system in detail.

TAGTIVITY: ACTIVITY TAGGING PROTOTYPE

We designed and implemented a prototype system called TAGtivity, which enables the user to easily assign a tag to any resource in their PC environment. The system generates comprehensive metadata about the created tags and a detailed log of the user's interactions with the system. This enables gathering quantitative data to aid the analysis of users' tagging practices.

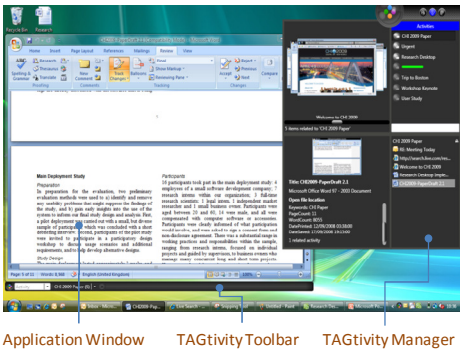

Figure 1. View of the TAGtivity prototype, showing the TAGtivity Manager (on the right) and Toolbar (below the application window).

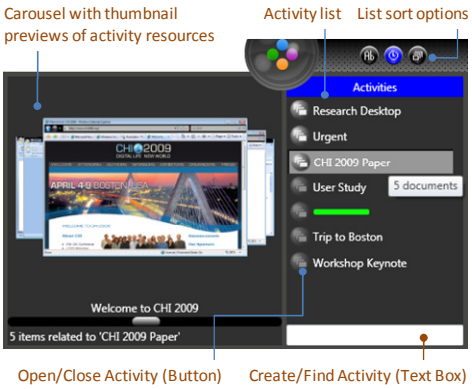

Figure 2. TAGtivity Manager, showing the list of tags and the expanded carousel window, which previews and enables launching resources linked with the selected activity.

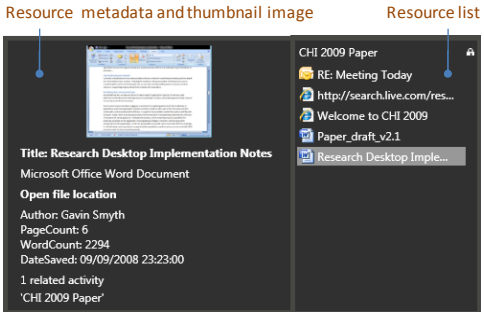

Figure 3. List of resources (right) and expanded overview window, showing thumbnail image and metadata for selected item (left). The resource list window slides-down from the top TAGtivity Manager window, when a tag is selected.

We anticipated that some tags would be used to designate tasks or activities that the user is performing. Thus, based on the review of prior work discussed above, we ensured that TAGtivity features are sufficiently general to support flexible gathering of resources in users' tasks.

TAGtivity comprises two UI components, the TAGtivity Manager and the TAGtivity Toolbar, which facilitate the creation and management of tags and tagged resources. Italso comprises a database store to persist information

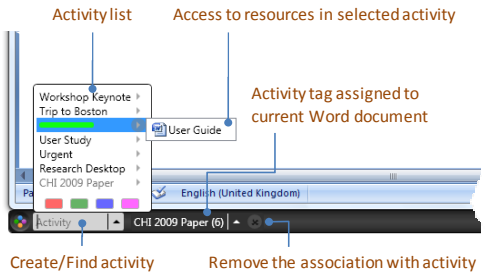

Figure 4. TAGtivity Toolbar, showing 1 tag associated with the open document and the expanded tag list. Multiple tags can be assigned to the document, by adding a new tag or selecting one from the list. Related resources can be accessed through a per-tag pop-up list.

about tags and their associated resources. The system is compatible with Windows Vista and XP operating systems and the Microsoft Office 2007 suite.

TAGtivity Manager

The TAGtivity Manager (TM), shown in Figure 1, is a centralized place for users to manage their activities and resources. It permanently displays a list of the user's tags, unless the user decides to close the display. By selecting one of the three buttons above the tag list, the user can sort the tag list alphabetically, by recency of use, or by group size (i.e., the number of associated resources). On mouse hover over a tag, the TM presents a sliding pane to the left (Figures 1-2) with a carousel of thumbnails and metadata to provide information and facilitate access to the resources associated with the tag. By clicking on a tag, a vertical pane slides down showing the list of associated resources in order of recency of access. On mouse hover over a resource, the horizontal pane provides more comprehensive metadata and a thumbnail image of the resource (Figures 1 and 3). By clicking on a resource name, the resource is opened in its default application. A right-click menu provides options for removing the resource from the list, i.e., disassociating it from the tag.

The TM supports a range of tag management functions. The text box allows the user to access a specific tag or to create a new one. By typing text into the text box the list of tags is filtered to show only matching tags. If the keyword is completely new, the user can select to use it as a new tag. Furthermore, by right clicking on a tag, the user can access options for deleting and renaming the tag.

TAGtivity Toolbar

In addition to the centralized management of user tags through the TM, we designed and implemented a TAGtivity Toolbar as an extension of the main MS Office 2007 applications: Word, Excel, PowerPoint and Outlook, and Internet Explorer 7 (IE7). Within the IE7 browser, each browser tab is handled independently.

TAGtivity Toolbars are located at the bottom of each application window (Figure 4). Using the text box on the left, the user can type a keyword to find existing tags or

create new ones. On mouse hover, a vertical pane slides up, showing the list of tags in reverse recency order. The user can attach a tag to the current resource by selecting a tag from the list or by typing in the text box. From here, the user can also tag a resource with one of four colours instead of using textual tags. Associated tags appear on the toolbar, showing in brackets the count of associated resources

From the list of tags displayed in the vertical slide pane, the user can click on the right arrow to view the respective lists of associated resources. The user can switch to or open a resource by clicking on the resource name. TAGtivity also enables users to associate files and folders with tags. The user can drag and drop an entire folder from the Desktop or from Windows Explorer onto the TM window to create another tag with the name of the folder or to associate the folder with an existing tag. Integration with the Windows Explorer also includes a context menu, which displays the tags associated with a resource. The drag-and-drop feature is particularly useful for expanding the tagging function to all file types, including resources that cannot be viewed in the Office 2007 applications. For example, the user can drop a PDF file onto a tag in the TM to associate it with that activity.

EVALUATION

As discussed above, our principle objective in deploying TAGtivity was to better understand the interplay between tagging and user behaviour in the context of information and activity management. In particular, we had the following broad research questions:

- What leads people to tag their resources?
- What aspects of resource and activity management do people perform through the use of tags?
- What impact does TAGtivity have on existing information management behaviour?

As our research questions are primarily related to organic work practices and emergent activities, we sought to investigate them by an in-situ study. We deployed TAGtivity to study participants over a period of 3 weeks, during which we observed their developing usage patterns. In the following sections, we first present the study methodology in more detail, and then the methods used for analysis.

Deployment Study

Preparation

We preceded our deployment study with two preliminary evaluation phases with the aim a) to identify and resolve any usability problems that might impact on the study findings, and b) to gain early insights into system usage in order to inform the design of the study methodology.

To address the first goal, we carried out a pilot study with 7 participants, which concluded with a short interview. Towards the second goal, we organised a participatory design workshop focusing on usage scenarios that emerged during the pilot study, additional requirements, and alternative designs for problematic system components or interactions. We then refined TAGtivity based on the collected feedback and suggestions.

Study Design

We deployed TAGtivity for 3 weeks. During this period, we conducted 4 interviews with each participant. Interviews were carefully designed to capture detailed information about, 1) participants' existing data management practices, 2) their use of TAGtivity and how this intersected with established practices and 3) how their tags mapped onto tasks and activities undertaken during the deployment. Data collection commenced with a pre-deployment interview (30 minutes), gathering demographic data and information regarding participants' roles, tasks, and current working practices. Following this, TAGtivity was installed and participants were given a one-on-one tutorial covering system features and usage. Participants also received a user manual. At the end of the first week, a 10-15 minute telephone interview was conducted primarily to address any questions or concerns that might be hindering participants' natural use of the system.

Two, in-depth interviews, totalling around 3 hours per participant, were conducted at the end of the second and third weeks with the aim to answer our research questions. The bulk of the analysis presented later on is based on the data gathered during these interviews. The second-week interview focused on the tags users created, the reasons for their creation, and the ways in which they were used. By comparing this motivation and usage with their existing practices, we were able to ascertain whether, and in what ways, information and activity management practices were affected. During the interview, automatically captured screenshots of significant events in the logs (such as when creating a tag or tagging a resource) were used as memory prompts to help participants recall reasons for their actions. The final interview focused more broadly on how the use of TAGtivity related to the structure of participants' responsibilities and work tasks. We also gathered information on situations when TAGtivity was found most useful, when tags were used less than expected, and when the software was decidedly not used. To aid this discussion, card sorting and participant-drawn diagrams were used.

Materials and Tools

We employed various data gathering techniques, most of which were built into TAGtivity. Usage logs collected detailed information (e.g., time, resource ID, tag ID, screenshot) about relevant user actions, such as the creation of a tag, or the tagging of a resource. The screenshots, logs and a summary of the user's current tag associations were gathered through an email feedback mechanism that was built into the TAGtivity system. The interviews were either audio or video recorded and were then fully transcribed. Finally, photographic evidence was collected of any physical material produced from the card-sorting and diagram-drawing sessions.

Participants

16 participants took part in the study: 4 employees of a small software development company; 7 research interns; 3 full-time research scientists; 1 legal intern, 1 independent market researcher and 1 small business owner. Participants were aged between 20 and 60, 14 were male, and all were compensated with computer software or accessories. Participants' working practices and responsibilities covered a broad range from interns with focus on a single project guided by a supervisor to business owners who manage many concurrent long and short term projects. Common to all the participants is that they made extensive use of their computers in performing their day-to-day tasks.

Analysis Methods

We employed three different methods to analyse the collected material: log analysis, profile generation, and undirected inductive coding of the interviews. Each is described below. Appropriate parts of the coding scheme are presented in tables in the findings section.

Log Analysis

The usage logs were analysed both for specific instances of events and for identifying usage patterns, such as tagging resources. Results of the log analysis are used in support of the findings discussed in the next section.

Profile Generation

We derived a profile for each participant, complete with ethnographic and work-role descriptions, summative information about their tags and associated resources, system usage statistics, and photographic evidence of materials generated during the final interview. These profiles were used both for reference and for discussion of the usage scenarios during our analysis.

Undirected Inductive Coding

In total, over 50 hours of semi-structured interviews were collected, and an undirected inductive coding method was used to formally analyse their content. The process was undirected as we did not begin coding with an existing model, but allowed a taxonomy to emerge organically from the process. According to the inductive approach prescribed by Thomas [16], initial codes were generated by multiple evaluators and from three data rich sample interview transcripts. The set of codes generated were analysed, categorized, merged, and reduced corporately by the team after processing each of the three sample transcripts until a stable coding scheme was agreed upon. The final coding scheme was validated by an independent assessor using Cohen's Kappa (κ=0.86). The remaining transcripts were coded by a single evaluator, using the final coding scheme.

STUDY FINDINGS

Over the 3 weeks, the 16 study participants created a total of 131 tags, 8.2(±5.9) tags on average per participant, and tagged 742 resources (average of 6.2(±5.2) resources per tag). Overall, TAGtivity was used 608 times to access a previously tagged resource. Figure 5 shows a scatter plot of the number of tags and resources each participant used.

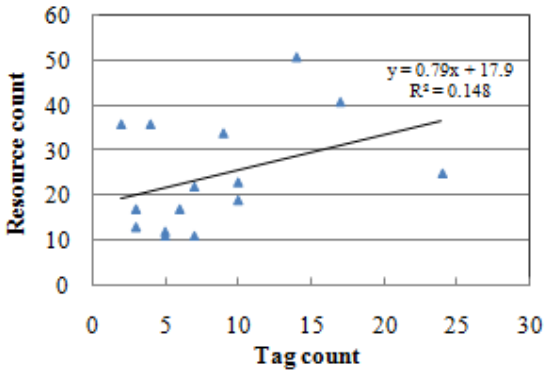

Figure 5. Scatter graph of the number of tags and resources for each participant (removed outlier corresponding to a user who created 5 tags and associated 119 resources)

Most users had 5 to 10 tags, with 2 or 3 items in each, during the study period. Notably, as the number of tags increases, the number of items associated with each usually remained low. One outlying case, containing a high number of resources was removed from the plot to enable a better view of the remaining points.

TAGtivity also supports non-textual colour tags. Of the 131 tags, however, only 10 were colour tags, 3 of which were immediately deleted after creation, 3 were named within 5 minutes, and one was renamed thereafter. Of the 121 textual tags, only 3 were renamed (1 immediately). These findings suggest that tagging is sufficiently simple and light-weight that people rarely needed intermediate grouping methods.

Tag Creation

In the following we structure our findings based on the notion of tag lifecycle, covering the creation and usage of tags. We begin with a categorization of the events that acted as triggers for tag creation. From our analysis of the interviews, we discovered four key triggers, each discussed in detail below: Place Holding, New Project, Tipping Point, and Time Saving. Definitions and examples of these triggers are provided in Table 1.

Place Holding. Eleven users created tags as place holders for future activities. For example, Nate created a tag to facilitate the gathering of interesting papers or links pertaining to robotics. Significantly, however, he did not add any items to the tag until a week later. Although the tag was created to gather resources, it was in fact triggered, not by any particular document or resource, but rather in anticipation of discovering resources at some future point.

Place holding also encompassed another behaviour, which refers to creating a tag in order to allow for the delayed handling of an activity related to an interruption, i.e. creating a type of 'to follow up' activity.

New Project. The creation of tags was often triggered by the onset of new projects. This behaviour was observed with 11 users. Unlike place holding, the tag is created with

Table 1. Tag creation triggers

Tag Trigger	Definition	Example	Direction
Place Holding	Tag is created with the expectation that it will be added to at a future point.	[Nate]: *I already knew that it will be not the main part of my research but if I find something, then it will be interesting to talk with my supervisor about it. This* [tag] *was mainly maybe for links I found or maybe papers in the SEM library that I found interesting.*	Forward looking
New Project	Tag is created at the outset of a project.	[Ruben]: *So, in the case of [tag], I was just starting to work on the project for the very first time, so I was about to review a specification and then do some development and interact with the user in checking some questions.*	Synchronous
Tipping Point	Tag is created at the point when sufficient resources have been gathered to warrant tagging.	[Ben]: *I had now gathered sufficient emails and sufficient files and sufficient work for us to want to start associating them together.*	Backwards looking
Time Saving	Tag is created after something took time to find and wanted to avoid doing so again if needed in the future	[Lois]: *When I found them after like 20 minutes looking through my folders, I actually added them to a category under the project name.*	Backwards looking

the intention of working on the activity right away, rather than creating a space to return to later.

Tipping Point. Another trigger for tag creation observed amongst 5 users is exemplified through a quote from Ben (see Table 1). In most instances, multiple resources pertaining to a current task had been assembled, and it was this assemblage which triggered the creation of a tag with which to group them. This is referred to as a tipping point, because before this the need to create the tag was not experienced, and after this point, resource organisation would be more costly or problematic.

Time saving. The creation of tags was also motivated when a long time had been spent locating a resource, and users wished to circumvent this process in the future. This behaviour was observed with 10 users.

In summary, the tag creation triggers discussed above have different temporal relationships with their current or future associated resources (see the last column in Table 1). In the case of Place Holding, the tag precedes the resource gathering phase. Place Holding tags may be the result of a deliberated plan or an unexpected interruption. In both cases, the triggers are *forward looking*, i.e. creating an activity that may be populated in the future. In the case of New Project, the creation of a tag is *synchronous* with the start of the work on the task, and gathering related resources. Tipping Point, in contrast, is a *backwards looking* strategy. The user has already gathered a set of resources, and the tag is used to impose a retrospective order or grouping upon these. Time saving is also backward looking, as the resources have already been located, although they are tagged with the aim to reduce the cost of re-visitation.

There was a significant preference for creating tags in context through the toolbar (75%, $t(30)=3.18$, $p<0.005$) rather than through the TAGtivity Manager, indicating that the majority of tags were created with a resource in hand. Ten tags, of the 131, were created but never associated with a resource. Furthermore, 35% of tags were created, often with a small number of resources, indicating that these tags were either created with a forward looking perspective or their usefulness was short lived.

Tag Usage

In this section we consider additional aspects of the *tag life cycle*: how they are used and how this use changes over time. We explore how the tags support the users' workflow, and complement and co-exist alongside their existing practices of data storage and access. Through this, we show how TAGtivity affords new forms of access to data, enabling new and useful groupings, which otherwise would not have occurred. We begin by charting general patterns of tag usage.

As mentioned before, 742 resources were tagged during the study. Of these, 100 resources were removed from the overall set of tags throughout the study. Most of these, however, occurred mostly in outlying examples (60 files were tagged unintentionally, by drag and drop of a folder onto TAGtivity Manager, and removed by the user soon after). Overall, on average 0.06 (±0.07) resources were removed from each tag.

Of the 608 times that resources were accessed (opened or switched to) using TAGtivity, 91.9% occurred through the TAGtivity Manager. Interviews revealed, however, that most participants were unaware that they could access related resources through the TAGtivity toolbar, thus the low usage of this feature, totalling only 49 accesses. Resources were accessed 415 times (68%) through the TAGtivity Manager's resource list (Figure 3), and 144 times (24%) through the thumbnail carousel (Figure 2), showing a significant preference for the former ($t(30)=2.22$, $p<0.05$).

Regarding the nature of the resources that were revisited through TAGtivity, the log analysis revealed wide diversity, indicating that users found tagging sufficiently flexible for many different resource types. Figure 6 breaks-down the resource accesses via TAGtivity by resource type. Overall, TAGtivity was widely used to revisit tagged e-mails (on 157 occasions), Web pages (on 98 occasions) and Word or PDF documents (on 174 occasions). These finding are significant as they confirm the importance of providing a

Table 2. Tag usage scenarios

Tag Usage	Definition	Example	Tag reach
Short term or transient	For an activity which is in a pre-organised state.	[Isaac]: *After a week, two weeks, I'm getting more of an idea of whether [this tag] is something that I'm going to want to keep, I'm going to want to create a sub folder, or it's just something I'm working on now but then I won't be.*	Single activity
Central repository	To collect resources from multiple sources (files, emails, etc.) into a single point, for easy access and activity resumption.	[Eric]: *I've been using it as a layer on top of my hierarchical directory structure, as a flat layer to keep track of multiple files that currently belong in different places, in one place.*	Mainly single activities
Filtering	To access key resources from a larger collection.	[Demitry]: *Yes, but then again, using TAGtivity to filter the most recent emails, the ones that are relevant. I found that very useful.*	Single and across activities
Meta-structures	To add a new organisation on top of file system structure.	[Lois]: *For example, I created a [tag...]PowerPoint presentation, those are core presentations of several clients, [...] And I also created by date, so I have 2007 and 2008, so some of the files that I created earlier and I can use it for a current project, I can easily access by year.*	Across activities

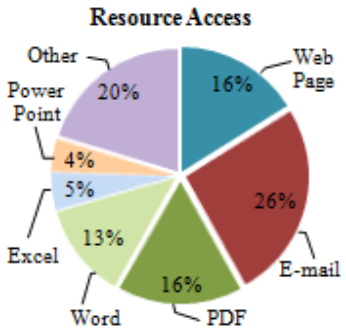

Figure 6: Comparison of the use of TAGtivity for accessing different types of resources, which cut across data silos.

unified approach for resource access, given that activities often require handling resources from multiple data silos.

The coding of the interviews revealed 4 main tag usage scenarios: supporting short term/transient activities; supporting revisitation of resources and activity resumption; filtering resources; and meta-organization of resources. These are described in Table 2 and are discussed next.

Support for short term/transient activities. An area where TAGtivity was found particularly effective was with short term tasks, or at the early stage of activities that later became persistent (this was true for 12 users).

The reason TAGtivity was valuable here was because it provided users with a method of associating resources, even before the task had reached a 'formal' level. For example, Isaac was using TAGtivity as a space to hold and group resources before their long term relevance became clear. For a short duration task, TAGtivity allowed him to manage the task up to its completion, at which point he removed the tag. However, for work of long term relevance, the use of TAGtivity was complemented with saving resources to his file system. Indeed, Isaac stated that one of the benefits he has derived from TAGtivity is that it has enabled him to delay creating folders, thereby creating fewer folders which he "wouldn't actually need" in the long term. Paul, a small business owner, suggested that, in his work, he considered tags to be an appropriate grouping mechanism for between 4 and (up to) 15 associated documents. For few resources, the overhead of creating a tag was too great, and for many resources he would favour organizing them within his current system.

One of the main findings of our study was that participants did not significantly alter their existing storage practices while using TAGtivity. Although participants created new groupings outside of, or prior to, existing storage practice, ultimately the storage practice was largely unchanged. Alongside this, we identified a number of usages based upon existing storage, where participants used TAGtivity to facilitate easy access to resources, filter current storage systems, or to create a meta-organisation of resources from the file system. Such usages are detailed next.

Central Repositories. The fact that TAGtivity enabled participants to gather resources across different data silos was universally expressed as a benefit (11 users explicitly mentioned this). In fact, participants often used TAGtivity specifically for this purpose, to create a single point of access to resources from multiple data silos. For example, Isaac used tags with the same name as the folders, which recurred throughout his storage system, in Outlook folders, and IE Favourites. TAGtivity was being used to funnel these distributed folders into one easily accessible place, which lends support to our initial conjecture that prior tagging systems' support for, typically, only one data silo did indeed act as a constraint on user behaviour.

Filtering. Another recorded use of TAGtivity was to filter key resources within a storage space (this usage was mentioned by 7 participants). The example quote given in Table 2 describes a situation where a large amount of email communication was being received over time. TAGtivity was found useful for fast access to the latest email at any one time. This was also true for fast access to the most essential files within a large folder.

Meta-level organization. Lois, an independent market researcher, used TAGtivity as a means to impose multiple views upon data contained within her existing storage structure. Multiple tags were assigned to resources to facilitate multiple access points to the underlying content. While Lois stored her data within individual project folders, this was not always useful for regular access, and she

frequently experienced problems remembering the precise location of resources. Her tags referred to individual project, type of project, document type, version, and year. Lois stated that TAGtivity helped her "organise my files without creating them, so it helped me group them, based on my processes and my needs".

3 participants used tagging for this meta-level organisation, while multi-tagging was widely observed across participants, with one exception of one participant. The maximum number of tags each participant applied to any given resource was on average 3.8 (± 4.1) (8 participants applied at most 2 tags per resource, 5 applied at most 3 to 5 tags and 2 applied more than 5 tags to at least one resource). This finding substantiates the importance of not constraining resources to belong to a single activity.

The tag usage scenarios considered above relate in different ways to participants' tasks. Tags used for short term project and central repositories tended to relate to just one task (see the Tag reach column in Table 2). However, tags used for filtering and meta-organisation typically supported multiple activities, especially in the latter case. We elaborate on this interesting finding in the discussion section below.

DISCUSSION

The findings of our study revealed a variety of ways that personal resource tagging can be of value to users of the PC environment. We can broadly group them into (1) emerging practices around management of transient and short term tasks, (2) high visibility and easy access of resources, and (3) diverse use of the generic tagging mechanism.

Transient and short term tasks

The TAGtivity system was found to support well transient and short term tasks as well as early stages of tasks when it is still unclear what the scope may be and what resources they may include. In such circumstances TAGtivity enables users to create holding spaces. These spaces sometimes act as an intermediary step before the resources are formally included into the organization structure, e.g., by creating a folder in the file system hierarchy. At other times, these holding spaces remain the only form in which these resources are associated together.

TAGtivity was specifically designed to afford lightweight tagging in context of desktop applications. Participants often commented on the ease and low overhead of creating tags. This perception appeared to facilitate the transient, early-stage, or forward looking resource collections. However, the nature of tagging also lent itself to the creation of transient collections. Tags act as pointers to content (rather than containers of it), an aspect of tagging which was widely understood in the sample. Deleting a tag is therefore of similar low overhead to creating one since this act only removes the pointer and not the content itself (in contrast to folders). It is the interplay between the TAGtivity system, the participants' patterns of use around folders and the specific affordances of tagging which we argue gave rise to the type of transient, forward looking and early-term tagging behaviours that we have observed.

Support for activity management

While TAGtivity is not an activity management application, it proved to support users in the execution of their everyday tasks. By enabling tagging from the application, TAGtivity allowed users to maintain multiple working contexts, as well as tagging resources to multiple tasks from a single point. This also supported interruption, enabling users to create placeholder tags without having to change their context.

Visibility of resources

Through the process of tagging and displaying resources at the application and desktop level, TAGtivity was used to surface information from the applications' data stores, such as e-mail and bookmark folders, and the file system organization. This applied generally across usage scenarios, but was especially central to the behavior of filtering, where participants specifically exposed important resources from the folders. By assigning a tag they surfaced them through the TAGtivity interface thus keeping them in view and within easy access. This served to raise awareness of particular resources, their importance relevant to other resources in the file hierarchy, and tasks related to them which needed to be attended to.

Generality of the tagging application

We found a great diversity in the ways in which users tagged resources. We observed users tagging by task, by type of task, by resources which cut across tasks, and by high level organizational strategy. Alongside this, users sometimes applied multiple tags to resources, organizing them along several dimensions at once. This diversity, we suggest, was a result of the open nature of TAGtivity, which was not focused or 'optimized' toward a specific area. We feel that this generality proved to be one of the foremost strengths of the system.

Design recommendations and future work

Based on these findings, we offer the following design suggestions for resource organization systems:

- Users should be offered intermediate workspaces explicitly. Our study spotlighted this as a currently unsupported area, absent from traditional systems.
- Resources should be surfaced from folder hierarchies and applications in a way that reflects the meaning of their associations.
- Future systems to support resource organization should enable greater flexibility.

From our findings, we identified three major areas to focus on for future redesign and investigation. First, it was clear that one reason why tagging was favorable was because it allowed links to be created across data silos to resources in their respective storage. Many participants reported that if they moved resources that they had previously tagged, the associations were lost. We are now altering the architecture of the software to allow tags to be stored with the files in

the filing system, so that this key benefit of TAGtivity is maintained. These changes will also permit sharing and collaboration of files and tags, which, although not studied in detail or reported here, were mentioned by many participants.

Our research also showed that colour tags served very little purpose, with some participants reporting that their meaning was easily lost. We seek to redesign the use of colour to augment tags, for facets such as importance.

Finally, our research revealed that although activities expired, there were very few instances of tag deletion. We seek to further investigate the requirements to support users in appropriately retiring tags, according to whether the activity resources need to be formalized, archived, or deleted.

CONCLUSIONS

In this paper we have described our investigations into how tagging in personal computing environments can be used effectively for both information and activity management. We (1) designed a generic tagging system, TAGtivity, that supports users in tagging and then managing any document in their personal computing environment; (2) performed an in-situ study of tagging practices with TAGtivity, using logging and periodic interviews; and (3) performed in-depth qualitative and quantitative analysis of the data collected. The study was in-situ, as opposed to benchmarking with other systems, as we sought to reveal insights into how tagging could support real-world activity and information management. Our findings revealed that TAGtivity supported users in both providing meta-organisations on top of their filing systems, and in supporting resource management within ephemeral activities that were previously unsupported.

ACKNOWLEDGMENTS

Thanks to Annika Hupfeld for her valuable contributions to this project and to the participants of the user studies.

REFERENCES

[1] Bannon, L., Cypher, S., Greenspan, S., and Monty, M.L. Evaluation and analysis of users' activity organisation. Proc. CHI'83, ACM Press (1983), 54-57.

[2] Bardram, J.E., Bunde-Pederson, J., and Soegaard, M. Support for activity-based computing in a person computing operating system. In Proc. CHI 2006, ACM Press (2006), 211-220.

[3] Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I. Taking email to task: the design and evaulation of a task management centered email tool. In Proc. CHI 2003, ACM Press (2003), 345-352.

[4] Cutrell, E., Robbins, D., Dumais, S., and Sarin, R. Fast, flexible filtering with phlat. In Proc. CHI 2006, ACM Press (2006), 261-270.

[5] Dourish, P., Edwards, W.K., LaMarca, A., and Salisbury, M. Presto: An Experimental Architecture for Fluid Interactive Document Spaces. ACM Trans. Comput.-Hum. Interact. 6,2 (1999), 133-161.

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

[7] Gonzalez, V.M. and Mark, G. "Constant, constant multi-tasking craziness": Managing multiple working spheres. Proc. CHI 2004, ACM Press (2004), 113-120.

[8] Henderson, D.A. and Card, S. Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. ACM Transactions on Graphics 5,3 (1986), 211-243.

[9] Hsieh, J.-L., Chen, C.-H., Lin, I.-W., and Sun, C.-T. A web-based tagging tool or organizing person documents on PCs. In CHI 2008 Workshop on Personal Information Management (2008).

[10] Jones, W., Phuwanartnurak, J., Gill, R., and Bruce, H. Don't take my folders away!: organizing personal information to get things done. In Proc. CHI 2005, ACM Press (2005), 1505-1508.

[11] MacIntyre, B., Mynatt, E.D., Voida, S., Hansen, K.M., Tullio, J., and Corso, G.M. Support for multitasking and background awareness using interactive peripheral displays. Proc. UIST 2001, ACM Press (2001), 41-50.

[12] Robertson, G., Horvitz, E., Czerwinski, M., Baudisch, P., Hutchings, D.R., Meyers, B., Robbins, D., and Smith, G. Scalable Fabric: flexible task management. In Proc. AVI 2004, ACM Press (2004), 85-89.

[13] Robertson, G., van Dantzich, M., Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., and Gorokhovsky, V. The Task Gallery: a 3D window manager. Proc. CHI 2000, ACM Press (2000), 494-501.

[14] Smith, G, Baudisch, P, Robertson, GG, Czerwinski, M., Meyers, B., Robbins, D., and Andrews, D. GroupBar: The TaskBar evolved. Proc. OZCHI'03 (2003).

[15] Tashman, C. WindowScape: a task oriented window manager. Proc. UIST 2006, ACM Press (2006), 77-80.

[16] Thomas, R. A General Inductive Approach for Analyzing Qualitative Evaluation Data. American Journal of Evaluation 27, 2 (2006), 237-246.

[17] Voida, S., Mynatt, E.D. & Edwards, W.K. Re-framing the desktop interface around the activities of knowledge work. Proc. UIST 2008, ACM Press (2008), 211-220.