

Lightweight Tagging Expands Information and Activity Management Practices

Gerard Oleksik[†], Max L. Wilson[#], Craig Tashman[‡], Eduarda Mendes Rodrigues^{*},
Gabiella Kazai^{*}, Gavin Smyth^{*}, Natasa Milic-Frayling^{*}, Rachel Jones[†]

[*] Microsoft Research 7 J J Thomson Avenue Cambridge CB3 0FB, UK {eduarda, gabkaz, gavin.smyth, natasamf} @microsoft.com	[†] Instrata Ltd. 12 Warkworth Street Cambridge CB1 1EG, UK {geoleksik, racheljones} @instrata.co.uk	[#] Future Interaction Technology Lab University of Swansea Swansea SA2 8PP, UK M.L.Wilson@swansea. ac.uk	[‡] Georgia Institute of Technology GVU Centre, Atlanta GA 30332, USA craig@cc.gatech.edu
---	---	---	--

ABSTRACT

Could people use tagging to manage day-to-day work in their personal computing environment? Could tagging be sufficiently generic and lightweight to support diverse ways of working and, perhaps, support new and efficient practices for managing applications and accessing documents? We investigate these issues by implementing the TAGtivity system that enables users to tag resources in the context of their ongoing work. We deployed TAGtivity and studied users' tagging practices in their actual work places over a three week period. Our analysis of interviews and logs reveals that affordances of the TAGtivity system supported users in a variety of information and activity management tasks. These include new practices for managing emerging activities and ephemeral information and accessing 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 the designs of systems that 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 encompasses a broad set of scenarios, including handling of multiple application windows, switching between tasks within and across applications, managing interruptions, and preserving the context of the current work. In contrast, information and file management 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 access to resources that the user requires during everyday work [2, 5]. Indeed, the user may need resources from potentially disparate parts of the file system. Furthermore, 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 (www.flickr.com) and Del.icio.us (www.delicious.com), have been studied extensively, we lack in-depth analysis of resource tagging within the PC environment.

Our work helps bridge this 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 a considerable effort in designing new and flexible tagging support but our primary objective was to observe and characterize the emerging tagging behaviors rather than evaluate the effectiveness of individual features or design options.

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

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

Copyright 2009 ACM 978-1-60558-246-7/09/04...\$5.00.

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.

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 rigid file organization. Furthermore, tagging enables capturing 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 focused 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 memorizing 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 high number of recurring folder names and organizational 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 behavior could be supported through better tools for cloning existing folder structures, their findings also suggests that an alternative organizational system which allows files to be organized along multiple, orthogonal dimensions at once could be useful. 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 (www.vistaglance.com), 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 focused 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 organize 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 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 focused 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 centralized activity bar which allows users to aggregate resources into groupings that can 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 is the 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 quantitative data aids the analysis of users' tagging practices.

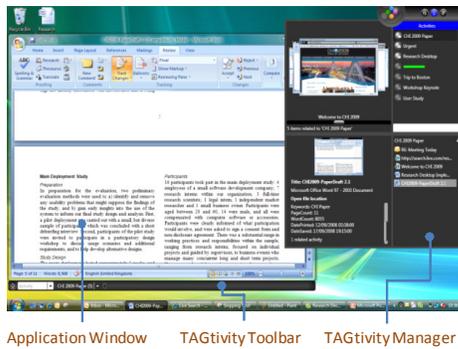


Figure 1. View of the TAGtivity prototype, showing the TAGtivity Manager (on the right) and Toolbar (below the application window).

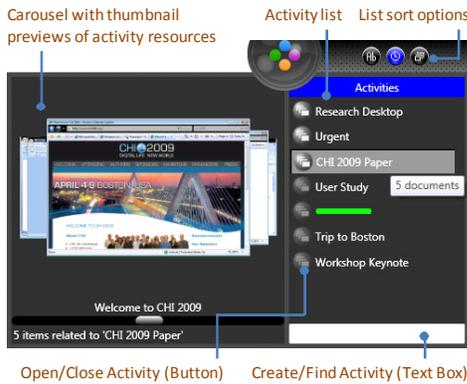


Figure 2. TAGtivity Manager, showing the list of tags and the expanded carousel window which provides previews and enables access to resources linked with the selected activity.

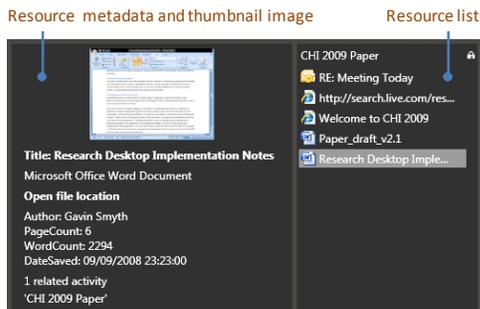


Figure 3. List of resources (right) and expanded overview window (left) showing a thumbnail image and metadata for a selected item. The resource list window slides-down from the top TAGtivity Manager window when a tag is selected.

We anticipated that some tags would be used to designate tasks or activities that the user is performing. Thus, guided by prior work (discussed above), we ensured that TAGtivity features are sufficiently general to support flexible gathering of resources during users’ tasks.

TAGtivity comprises two UI components, the TAGtivity Manager and the TAGtivity Toolbar, which facilitate the

creation and management of tags. It also includes a database store to persist information about the usage of tags and tagged 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 tags alphabetically, by recency of use, or by 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 and 2), with a carousel of thumbnails and metadata about the tagged resources. One can access a resource by clicking on the thumbnail image.

Clicking on a tag invokes a vertical pane that slides down showing the list of associated resources in the order of recency of access. On mouse hover over a resource, the horizontal pane provides metadata and a thumbnail image of the resource (Figures 1 and 3). By clicking on the name, the resource is opened in its default application. A right-click menu provides options for removing it 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 not found in the list, the user can choose to use it as a new tag. Furthermore, by right clicking on a tag, the user can access options for deleting and editing the tag.

TAGtivity Toolbar

In addition to 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).

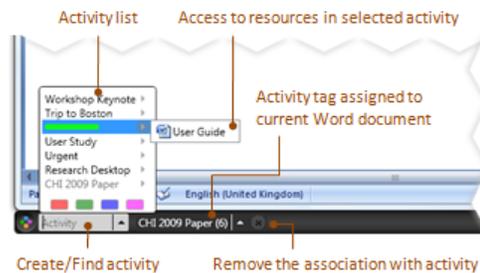


Figure 4. TAGtivity Toolbar, showing one tag associated with the open document and the expanded tag list. Multiple tags can be assigned to the document by typing in a new tag or selecting one from the list. Related resources can be accessed through the list associated with a tag.

Using the text box on the left, the user can type a keyword to find existing tags or create new ones. Associated tags appear on the toolbar, showing in brackets the count of associated resources. On mouse hover, a vertical pane slides up, showing the list of tags in reverse order of recency. The user can attach a tag to the current resource by selecting a tag from the list or by typing in the text box. Alternatively, the user can attach color tags instead of textual ones.

From the list of tags displayed in the vertical slide pane, the user can click on the right arrow to view the lists of associated resources. The user can open or switch to a resource by clicking on the resource name.

Integration with the File System. TAGtivity also enables the user 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. If it is dropped onto an existing tag, the folder or its constituent files are associated with the tag. Otherwise a new tag is created with the name of the folder. Integration with the Windows Explorer further includes a context menu associated with files and folders that displays the associated tags.

The drag-and-drop feature is particularly useful for expanding the tagging function to all file types, including those that cannot be viewed in the MS Office 2007 applications and thus cannot be tagged through the TAGtivity Toolbar. For example, the user can drop a PDF file onto a tag in the TM to associate it with that activity.

USER STUDY

The main objective of our research is to understand the interplay between tagging and user behavior in the context of information and activity management. In particular, we would like to answer the following research questions:

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

As these are related to organic work practices 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.

TAGtivity Deployment

Preparation

We preceded the in-situ 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 the system usage in order to inform the design of the study methodology.

To meet the first goal, we carried out a pilot study with 7 participants, which concluded with a short interview.

Towards the second goal, we organized a participatory design workshop that focused on the usage scenarios that emerged during the pilot study. This led us to consider additional requirements and explore alternative designs for system components or interactions. We refined TAGtivity based on the collected feedback and suggestions.

Study Design

We deployed TAGtivity for 3 weeks. During that period we conducted 4 interviews with each participant. Interviews were carefully designed to capture detailed information about: 1) participants' existing data management practices, 2) use of TAGtivity and effects that has on established practices and 3) relations of tags to projects, tasks, and activities that participants conducted during the time of the study.

We started with a pre-deployment interview (30 minutes), to gather demographic data and information about participants' roles and responsibilities, tasks, and work practices. Following that, we installed TAGtivity and conducted one-on-one tutorials covering system features and functions. We also provided a user manual. At the end of the first week, we followed up with 10-15 minute telephone interviews, primarily to address questions or issues that might be hindering participants' use of the system. At end of the second and the third weeks we conducted two in-depth interviews, respectively, lasting about 3 hours in total per participant. The main objective was to answer our research questions. In fact, the bulk of analysis we present later in the paper is based on the data gathered during these interviews.

In the second-week interview we focused on the tags that users created, the reasons for creating them, and the ways they were used. By comparing the motivation and usage of tags with their existing work practices, we were able to ascertain whether and in what ways the user's behavior was affected. During the interview, we used information from the logs to help participants remember their actions. In particular, we designed our logging software to capture screenshots of significant events such as creating a tag, assigning a tag to a resource, etc., and use them as memory prompts.

In the final interview we explored how the use of TAGtivity related to the structure of participants' activities and projects. We gathered the participants' feedback on situations when TAGtivity was found most useful, when used less than expected, and not used at all. To aid this discussion we used card sorting and diagram drawing, asking participants to depict the scope and structure of their work.

All the interviews were audio or video recorded and then fully transcribed. We also took photos of materials produced during the card-sorting and diagram-drawing sessions. This rich record of user data enabled us to conduct qualitative and quantitative analyses of the user experience with TAGtivity.

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 male and 2 female. For their participation, they were compensated with computer software or hardware accessories.

Participants' working practices and responsibilities ranged from interns who focused on a single major project, guided by a supervisor, to business owners who managed multiple concurrent projects, involving a short-term or long-term effort. All the participants extensively used computers in their day-to-day tasks.

Analysis Methods

Processing of collected data involved three methods: analysis of the user logs, generation of user profiles, and undirected inductive coding of the interviews. Each is described below.

Log Analysis

We analyzed usage logs for specific system events and overall tagging patterns. We used the log analysis in combination with other data to provide evidence for our findings, discussed in the following section.

Profile Generation

For each participant we derived a profile that comprises a demographic and work-role description, summative information about tags and associated resources, system usage statistics, and materials generated during the final interview – card sorting and diagram drafts. These profiles were used both for reference and for discussion of the usage scenarios during our analysis.

Undirected Inductive Coding

In total, we collected over 50 hours of data from the semi-structured interviews and applied an undirected inductive coding method to analyze their content. We did not begin coding with an existing model but allowed a code taxonomy to emerge organically from the process.

In accordance with the inductive approach prescribed by Thomas [16], initial codes were generated by multiple researchers from a sample of three interview transcripts. The researchers processed each of the three transcripts, adding new codes when needed to describe new aspects found in the users' statements. This process was completed when the coding scheme became stable, i.e., when no new terms were found worth adding. The generated codes were analyzed, categorized, merged, and reduced and the final coding scheme validated by an independent assessor using Cohen's Kappa ($\kappa=0.86$). One of the evaluators then applied the codes to the remaining transcripts, preparing them for the qualitative and quantitative analysis. We show some parts of the coding scheme in Tables 1 and 2.

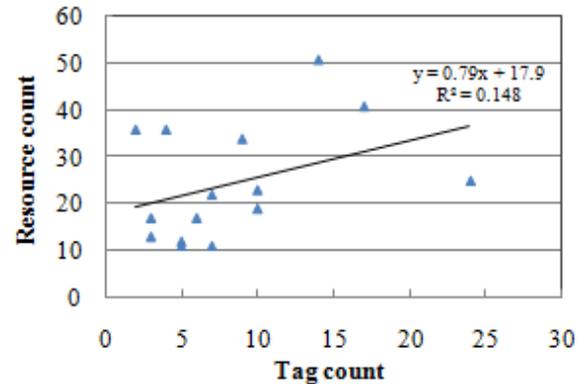


Figure 5. Scatter plot of the number of tags and resources for each participant, without the outlier that corresponds to a user who created 5 tags and associated 119 resources to them.

STUDY FINDINGS

Over 3 weeks, 16 study participants created a total of 131 tags, i.e., 8.2 (± 5.9) tags on average per participant, and tagged 742 resources, i.e., assigned 6.2 (± 5.2) resources to a tag on average. They used TAGtivity 608 times to access a previously tagged resource.

Figure 5 shows a scatter plot of the number of tags and resources that each participant used. During the study period most of the users created 5 to 10 tags with 2 or 3 items per tag. 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 color tags. Of 131 tags, however, only 10 were color tags, 3 of which were immediately deleted after creation, 3 were replaced by a textual tag within 5 minutes, and one was renamed thereafter. Of 121 textual tags, only 3 were renamed (1 immediately). These findings suggest that textual tagging is already sufficiently simple and efficient that people rarely needed a more lightweight method for intermediate grouping of resources. In the following section we present our findings in detail describing the observed tag lifecycle and usage scenarios. We begin with a categorization of events that acted as triggers for tag creation.

Tag Creation

From the interview analysis, we discovered four key triggers for creating a tag: 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 intended to gather resources, the trigger was not any particular document or resource but rather the user's

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>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.</i>	Forward looking
New Project	Tag is created at the outset of a project.	[Ruben]: <i>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.</i>	Synchronous
Tipping Point	Tag is created at the point when sufficient resources have been gathered to warrant tagging.	[Ben]: <i>I had now gathered sufficient emails and sufficient files and sufficient work for us to want to start associating them together.</i>	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]: <i>When I found them after like 20 minutes looking through my folders, I actually added them to a category under the project name.</i>	Backwards looking

anticipation of discovering resources at some future point. Place holding also encompassed another behavior. The tags were created to allow for delayed handling of an activity, e.g., ‘to follow up’ with an action item after an interruption.

New Project. The creation of tags was often triggered by the onset of new projects. This behavior was observed with 11 users. Unlike place holding and creating a space to return to later, these tags were created with the intention to work on the activity right away.

Tipping Point. With 5 users we observed another tag creation trigger, i.e., the perceived value of the resources assembled through a task or activity. As exemplified through the quote from Ben (see Table 1), the user may have assembled multiple resources pertaining a current task and realized the importance and advantage of keeping them together for future reuse. This is referred to as a tipping point—prior to this point the creation of a tag was not considered. However, it was then realized that the cost of reassembling resources is significant and that having a tag that keeps them together would be beneficial.

Time Saving. Similarly, tags were motivated by the need to avoid time consuming processes. For example, a resource may have been discovered through search and browsing that took a considerable time. By bookmarking the item using a tag, the users circumvented the need to engage in the same process again. This behavior was observed with 10 users.

In effect, these tag triggers are characterized through their temporal relationships with associated resources (see the last column in Table 1). In the case of Place Holding, the tag precedes the resource gathering. The tag may be the result of a deliberate plan or an unexpected interruption. In either case, the triggers are *forward looking*, creating an activity space that may be populated with resources in the future. In the case of New Project, the creation of a tag is *synchronous* with the start of the work and use of resources. Tipping Point, in contrast, is *backward looking*; the user has already gathered a set of resources and the tag provides their retrospective order or grouping. Time Saving falls into the same category, with more emphasis on the amount time

used to access a resource than on the number of resources that need to be assembled.

From the usage logs we observe a significant preference for creating tags through the Toolbar, i.e., in the context of an application (75%, $t(30)=3.18$, $p<0.005$) rather than through the TAGtivity Manager. This suggests that the majority of tags were applied to a resource in hand; only 10 out of 131 tags were created but never associated with a resource. We note that 35% of tags had only a small number of associated resources, indicating that these tags may have been created with a *forward looking* perspective or were of short time use.

Tag Access

In this section we consider how tags are used and how their usage changes over time. We explore the ways they support the user’s workflow and co-exist alongside existing practices in data storage and access. Our findings show that TAGtivity affords new forms of access to data by enabling useful grouping of resources that otherwise would not have occurred. We begin our discussion by charting general patterns of tag usage.

Deleting Tags and Resources. As mentioned before, 742 resources were tagged during the study. Of these, users disassociated 100 resources from the tags they created. Most of these, however, occurred in the specific case that is considered an outlier: 60 files were tagged unintentionally by drag-and-drop of a folder onto the TAGtivity Manager and removed by the user soon after. Overall, on average 0.06 (± 0.07) resources were removed from each tag.

Accessing Tags and Resources. Out of 608 accesses to tagged resources, i.e., opening or switching to a resource, 91.9% occurred through the TAGtivity Manager. Interviews revealed that most participants were unaware that they could access individual resources through the TAGtivity toolbar by clicking on the arrow and expanding the resource list (Figure 4). This possibly contributed to the low usage of that feature, totaling only 49 accesses. Resources were accessed 415 times (68%) through the TAGtivity Manager’s resource list (Figure 3) and 144 times (24%)

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]: <i>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.</i>	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>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.</i>	Mainly single activities
Filtering	To access key resources from a larger collection.	[Demitry]: <i>Yes, but then again, using TAGtivity to filter the most recent emails, the ones that are relevant. I found that very useful.</i>	Single and across activities
Meta-structures	To add a new organisation on top of file system structure.	[Lois]: <i>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.</i>	Across activities

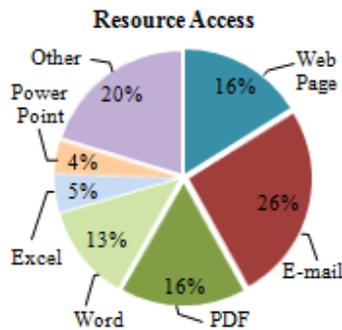


Figure 6: Comparison of the use of TAGtivity for accessing different types of resources, which cut across data silos.

through the thumbnail carousel (Figure 2), showing a significant preference for the former ($t(30)=2.22, p<0.05$).

The log analysis revealed a wide diversity of resources that were tagged and accessed using TAGtivity, suggesting that the system is sufficiently flexible to cover different resource types. Figure 6 provides a break-down of resource accesses by 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 findings confirm the importance of providing a unified approach to resource access, given that activities often involve multiple applications and proprietary data stores. TAGtivity can potentially help with crossing data silos and increasing the user productivity.

Tag Usage

The coding of the interviews revealed 4 main tag usage scenarios: managing short term/transient activities, revisiting resources and resuming activities, filtering resources, and creating meta-organization of resources. These are described in Table 2 and are discussed next.

Managing short term/transient activities. TAGtivity was found particularly effective for managing short term tasks and early stages of longer term activities. This was observed with 12 users. Indeed, TAGtivity provided users with the ability to collect and associate resources before a task is well formulated.

For example, Isaac was using TAGtivity to hold and group resources before their long term relevance became clear. Conveniently, for short term tasks, TAGtivity helped him manage resources up to the task completion, at which point he removed the tag. For work of long term relevance, Isaac complemented tagging with saving resources to his file system. He stated that TAGtivity provides a unique benefit of enabling him to postpone creation of folders and thus reduce the number of those that he “wouldn’t actually need” in the long term. Paul, a small business owner, suggested that, in his work, tags are an appropriate grouping mechanism for sets between 4 and 15 related resources. For very few resources, the overhead of creating a tag is too great and for a large number of resources he prefers to organize them within the current system.

One important finding of our study is that participants have not significantly altered their existing practices in storing and organizing files. TAGtivity is not replacing but rather complementing them by facilitating access to data across storage types, filtering key resources, and creating a meta-organization over the file system structure.

Central Repositories. The fact that TAGtivity enables access to resources across the file system and application specific data stores was universally expressed as a benefit by the participants. Some used TAGtivity specifically for the purpose of creating a single point of access to resources from multiple data silos (this usage was mentioned by 11 participants). For example, Isaac created tags with the same name as folders that already existed within his file system, in his e-mail (MS Outlook Folders), and among the browser bookmarks (IE Favorites). He used TAGtivity to funnel these distributed folders into one easily accessible virtual place. This further supports our initial conjecture that restricting a tagging system to support only one data type acts as a constraint on the user behavior.

Filtering. Another beneficial use of TAGtivity was to filter key resources within a storage space (this usage was mentioned by 7 participants). The example quote by Demitry, in Table 2, describes a situation where a large amount of email communication was received over time. TAGtivity enabled him to tag and keep visible the most

useful recent e-mails. He used the same practice to ensure fast access to most essential files within a large folder.

Meta-level organization. Lois, an independent market researcher, used TAGtivity to create multiple views upon the data stored within her file system. While she organized data in individual project folders, that structure was not always useful for regular access and she often had difficulties finding the precise location of resources. Lois eased this problem by using TAGtivity to assign multiple tags to individual resources and thus providing multiple access points to the data. Her tags reflected the project names, project types, document types, document versions, and the year. Lois stated that TAGtivity helped “*organize my files without creating them [folders]; so it helped me group them based on my processes and my needs*”. Similar practices were observed with two other participants.

Generally, multi-tagging was widely observed across participants, with the exception of one participant. Looking at the maximum number of tags that each individual applied to resources, we found that the average is 3.8 (± 4.1) tags. Eight participants used at most 2 tags, five participants used a maximum of 3 to 5 tags, and two participants used more than 5 tags for at least one resource. This confirms that we should not constrain resources to a single activity.

Our study also suggests that assignment of single or multiple tags may depend on the type of activity they refer to. Tags associated with short term projects and central repositories rarely apply to other tasks (see the Tag reach column in Table 2). In contrast, the tags used for filtering and meta-organization of resources typically support multiple activities and thus may apply to resources with other tags.

DISCUSSION

The study of TAGtivity revealed several distinct characteristics of tagging in the PC environment. Participants benefited from (1) support for managing transient and short term tasks, (2) increased visibility and easy access to relevant resources, and (3) diverse use of the generic tagging mechanism.

Transient and short term tasks. The TAGtivity system was found to support transient and short term activities as well as early stages of longer term tasks when the scope is still unclear. In such circumstances the tags enable users to hold together potentially relevant resources, often serving as an intermediary step before the resources are included into the storage and organizational structure, e.g., by creating a folder in the file system hierarchy. However, tags may also remain the only form in which these resources are associated together.

TAGtivity was specifically designed to afford lightweight tagging in the context of desktop applications. Participants often commented on the ease and low overhead of creating tags. This perception made tagging attractive even for the most transient activities. However, tagging also lent itself to

the creation of virtual groups and collections that could be easily dispersed when not further needed. Indeed, tags act as pointers to content rather than content containers. Deleting a tag, therefore, is low risk, almost non-consequential, since such an act only removes the pointers to the content and not the content itself. We believe that it is this fine interplay between file system and the affordances of the TAGtivity system that gave rise to the transient, forward looking and early-stage tagging behaviors that we have observed.

While TAGtivity is not an activity management application in the sense of [2, 12, 14], it has proven to support users in performing their everyday tasks. By enabling tagging from applications, TAGtivity provides a unique benefit of allowing users to maintain multiple working contexts and associate resources with multiple tasks without changing the current context. This also enables users to handle interruptions, e.g., by creating a placeholder tag and associate resources with the emerging activity without shifting the focus of their work. Alternatively, they can briefly switch to another work context by accessing a relevant resource directly from the current application.

Visibility of resources. Through the use of TM and Toolbar to display tags and resources at the application and desktop level, TAGtivity enables the user to surface information that is otherwise buried in application specific data stores or the file system hierarchy. This is central to the filtering practice where participants used tags specifically to expose important resources from file folders and e-mails. Increased visibility of resources has an important side effect. It raises the awareness of tasks related to the tagged resources and serves as a reminder of activities and issues.

Generality of the tagging application. We observed a great diversity in the way users tagged resources. They grouped them by task properties but also experimented with different strategies for achieving alternative resource organizations or unifying resources across data stores. Sometimes users applied multiple tags to resources in order to organize them along several dimensions at once. This diversity, we believe, is a result of the generic nature of TAGtivity—the system was not aimed or optimized for a particular usage scenario. This generality has proven to be one of the foremost strengths of TAGtivity.

Design recommendations and future work

Based on the study findings, we offer several recommendations for the design of resource and task management systems:

- Users should be offered intermediate workspaces that support various task stages. This is absent from traditional and contemporary systems.
- Resources should be surfaced from folder hierarchies and applications in ways that explicate their associations and reflect the relations to the user’s work.

- Systems that support resource organization and management should provide greater flexibility.

From the users' feedback on TAGtivity usage we identified three major areas for future redesign and investigation. First, many participants reported that if they moved a resource that they had previously tagged away from its original storage, the tag associations with the resource were lost. This is particularly acute with e-mail messages that are frequently archived and affects one of the important TAGtivity benefits—linking resources across application data silos. We are now altering the system architecture to enable tags to be stored directly within resources. Thus, the tags remain with the resource as it is moved around the storage locations.

The same change would enable tag sharing in collaborative scenarios. Although we have not elaborated on this aspect in the paper, several study participants suggested possible benefits from tag sharing.

The study also showed low use of color as a sole means of tagging. In future designs we will explore the effect of color as augmentation to textual or graphic tags.

Finally, we observed relatively rare instances of tag deletions, despite the fact that activities may have expired. Thus, we will explore ways to support users in the management of the complete tag lifecycle.

CONCLUSIONS

In this paper we describe our investigation of resource tagging in the personal computing environment and demonstrate that tagging can support novel ways of managing resources and user tasks. The research was facilitated by a generic tagging system, TAGtivity, that we designed to enable lightweight tagging of resources in the PC environment. We conducted an in-situ user study to observe emerging tagging practices through user logging and periodic interviews. We performed in-depth qualitative and quantitative analyses of the collected data. The study results show that generic tagging, as provided by TAGtivity, can support users in managing ephemeral activities, provide easy access to information from diverse data stores, and enable alternative meta-organizations of the file system to support dynamic and diverse ways of working.

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 personal 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 evaluation 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 for 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., Volda, 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., Ridsen, K., Thiel, D., and Gorokhovskiy, V. The Task Gallery: a 3D window manager. Proc. CHI 2000, ACM Press (2000), 494-501.
14. Smith, G., Baudisch, P., Robertson, G.G., 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. Volda, 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.