Visualization systems as navigational aids around the web of personal information : a survey

Mohan Raj Rajamanickam, mohanr@cs.ubc.ca

Department of Computer Science, University of British Columbia

Abstract

There are numerous tools to organize and visualize personal information, yet an ideal tool eludes many users. PIM (Personal Information Management) tools are probably one of the most abandoned class of software. But I believe that a right tool has the potential to make the most difference in the life of an average user. This survey succinctly reviews many such tools with a focus on those centered around the world wide web. The aim is to enlighten an enthusiastic developer, who wants to develop yet another PIM tool, of the challenges involved. The discussion is limited to a subset of the tools which use any form of visualization for conveying the information structure the user.

1. Introduction

This is the quest of a person searching desperately for a solution, battered by continuous interrupts and overwhelmed by tasks, to help him in his ever failing attempt to organize multitudes of lists, heaps of tasks, organize his work, and ultimately his life; trying to make sense of the madness and onerously trying to get stuff done all the while. The focus of this paper is to better understand the underlying

problem. The contribution of this paper is ideally two-fold: to have a better idea about the field and efforts made by other researchers which should 1. primarily be of help, by providing a rough set of guidelines, to anyone who is considering building yet another PIM, outliner, task scheduler, project manager, todo-list ..etc, which adds one more trickle into the ocean of existing tools, and additionally 2. be likely an overkill of information to an enthusiast who is looking out for a solution and wants to have a sense of the range of available solutions and why some might work or not work for his purpose. Specifically, I am interested in looking at visualization systems that would help me to store, organize, and display the information collected over the course of sense-making tasks (where Sense-making is defined as a *motivated, continuous effort to understand connections in order to .. act effectively*).

1.1 The Problem

Until a little while ago, people typically had more time and attention when compared to the availability of information that they were interested in - both in terms of quantity and quality. But the dawn of the information age has caused a reversal of the ratio. People now have access to more information than they can deal with effectively. This reversal of roles of abundance between attention and information creates new problems in terms of managing the information overload and necessitates building of tools to help us in the process of making sense of the torrent of information flowing our way and also help us figure out where/how to store the information that interests us in a way that facilitates easier retrieval. As a community we have solved the problem of locating information to a considerable extent. We now have to deal with the problem of organizing all the found information in a way that minimizes the cost of accessing it again when we need it as observed by Jones et al. (2001).

1.1.1 Wild Wild Web

The problem is acute in the realm of the world wide web. Kandogan and Shneiderman (1997) observed about more than a decade back that web browsers did not provide support structures to give the user a sense of overview and a sense of the current location in the complex information space being explored. Nor did they facilitate organizing and filtering of information or aid users in accessing previously visited pages without incurring a high cognitive load. I believe that the situation has not improved in the decade since and a considerable number of users are now plagued with the problem more severely than ever before with the web becoming a vital part of their core work-flow. Since most part of finding and managing information occurs in the web-browser, at least for a subset of users like me, I am specifically interested in reviewing systems that work with or as-part-of web-browsers helping users organize their information primarily around browser-based information artifacts such as tabs, bookmarks and browsing history.

Thomas W. Malone (1983) found that one of the primary problems that people had with information organization systems is the cognitive load incurred in classification, in determining where/how to file a certain piece of information. I think that having the personal information organization system based on web-based artifacts such as URLs lends itself well to this problem by

1. having an implicit relation based on user action, for example: this URL was opened from this URL and hence should be filed under it or associated with it etc.
2. the possibility of taking advantage of semantic structures of the information for automatic classification of information artifacts.

1.2 The Minimum effort rule

Working in a web-browser also has the inherent advantage that the system becomes relatively more

portable across computers. This is important as an increasing number of users are accessing and organizing information using multiple devices over the course of a day to accomplish their work and any information-organizing tool that is not available to the user when they want to collect a piece of information reduces the motivation for the user to continue using the tool [Bernstein et al., 2008]. Jones et al. (2001) find that users tend to capture information outside the tool even at the slightest inconvenience caused by systems whose affordances do not match their needs and expectations. Tool authors should bear in mind that users are more sensitive to the effort-put-in to value-derived ratio of meta tools, such as the information organization tools, which are not part of the core primary task of the user when compared to those tools that are essential to complete the task at hand. Moreover, when the user is in the *state of flow* (a state of mind at which creativeness and concentration are in sync and at their peak)*,* they would like the capture/retrieval of information to present as little overhead as possible in order to not disturb their *state of flow.* For every bit of information that is captured outside the system, because the system is either not available at that point of time or not flexible enough to accommodate the information at hand etc, the life-span of the system under the user gets reduced.


1.3 Web-browser - in or out ?

There are several ways in which a visualization based information organization system could be made to work with the browser. You could implement it as a browser plugin that runs in a tightly-coupled manner as part of the browser. The advantage of this approach is the range of data (low level user events for example) that you gain access to as the developer of the tool. The user also benefits from using a single extended application (for example InsightFinder [Cheng and Gotz, 2009] is an extension to the standard web browser that assists users in their sense-making tasks on the web). The main disadvantage with this kind of design being the inherent restrictions of making use of the web-browser as a powerful visualization platform. But this is becoming less of a problem with the availability of

visualization platforms in forms that can be natively used in the browser (for example as Javascript libraries) and the current crop of web-browsers moving towards GPU accelerated rendering of the content.

The other method is to implement the visualization as a system outside the browser and have it running in tandem with the browser. Mayer and Bederson (2001), for example make their system act as a proxy server which the users can point their browsers at. The system then captures the URLs visited by the user and uses a relatively powerful visualization tool to render the data. While an external tool is powerful in terms of the flexibility it offers to the developer and visualizations that could be built, the user has to use two different tools in tandem to accomplish the task – which doesn't go well with *the Minimum effort rule* that we discussed earlier. The system data, which is coarse grained than the earlier in-built approach, has to reside outside the browser and this makes it difficult to provide tight integration with existing information, for example bookmarks etc, that the user has accumulated in the browser.

2 Searching vs Browsing

So, can't this problem be solved by creating a personal search engine that indexed most of the information that a person ever came in contact with, like for example the *Stuff I've seen* system by Dumais et al. (2003) ? I believe that such an approach could help us solve the problem partially. To understand why search can't solve the problem completely, we need to understand search vs browse interfaces and how/when people use them.
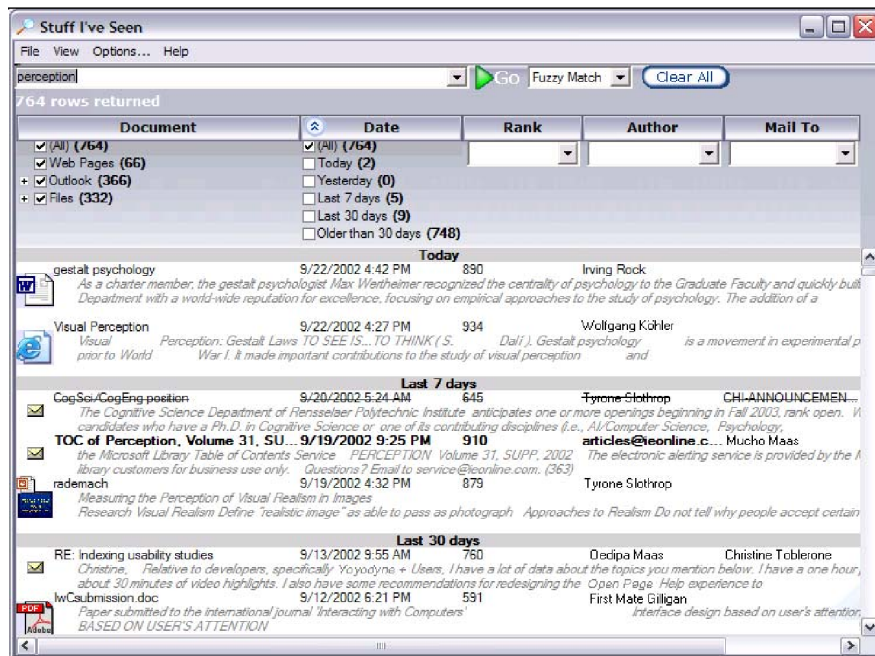
Stuff I've seen, a system to index personal information

Knowledge/information interaction interfaces can either be modelled as search-based or browse-based. Search-based interfaces are typically faster in giving access to certain information to the user, but requires the user to have prior knowledge and remember or come up with key terms. Browsing interfaces are typically slower, but have lower effective cognitive load on the user, in that they do not require the user to remember/recollect as much information as the search-based interfaces. Browsing interfaces have lower barrier of adoption and also help the user discover new information or information that had been forgotten/overlooked. But browsing interfaces typically fail to scale well with increase in density/volume of information. We shall pause the discussion briefly and return to it after discussing work-spheres.
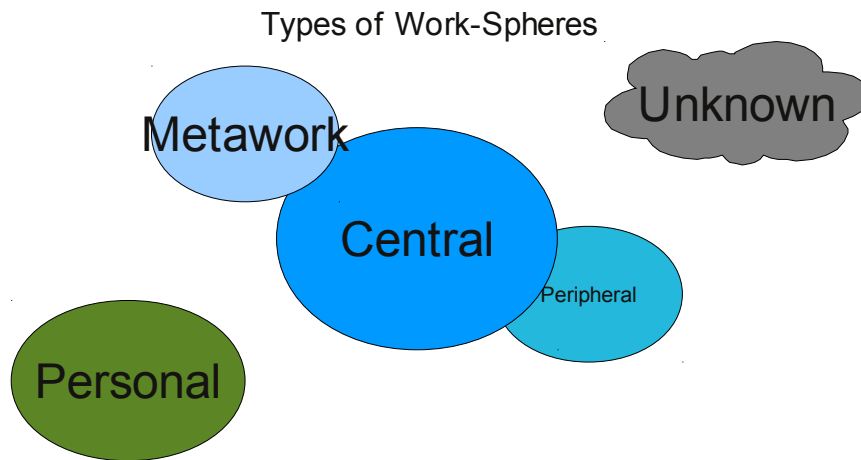
2.1 Work-spheres

González et al. (2004) did a field study about how people dealt with information in their workspaces. They found that participants organized their work implicitly at a higher level of abstraction in terms of work-spheres and effectively dealt with around ten work-spheres in the course of a day. On average participants worked on a sphere for roughly about 30 minutes with intermittent short interrupts, that may be self-initiated or external, that forces them to switch momentarily to other spheres or among different tasks within the same sphere.

Participants generally expressed desire to stay within the same sphere without interruptions until they made a satisfactory progress on the current task. But González notes that this rarely ever happens. Hence people use different artifacts such as sticky notes, emails to self, task lists, project management tools etc as coping mechanisms to help them maintain a sense of their tasks in different spheres and as a reminder to come back to the primary task after attending to an interrupt. González et al. (2004) find that the current tools available for information technology workers fail to take into consideration the different spheres and the constant, frequent switching among tasks. They hence fail to help the user in providing mechanisms for effective sense-making and recovery from interrupts.

González et al. characterize an effective artifact that aids the user in this scenario to be one that provides a two pronged function:

1. It should act as a '*signal*', as a reminder towards the working sphere that needs to be attended to,

2. It should '*describe*' in some detail as to what needs to be attended to.

Types of Work-Spheres

González et al. (2004)

2.2 Searching vs Browsing vs Work-spheres

For our purpose, the ideal interface would be browse-based, given the limited number of work-spheres that users typically deal with during the course of the day. But the interface should also be augmented with search interfaces considering the frequency and number of interrupts that occur during the work day. This would enable the user to have the best of the both worlds, rapidly switching between tasks using search interface when he remembers key terms and using the browse interface to explore and prime the information space, to get a sense of the outstanding tasks and to monitor progress etc.

2.3 Is it safer to assume that users are lazy ?

Builders of information organization tools often assume that people are willing to put in the additional effort in house keeping to reap the benefits of using the tool. But the reality is a bitter pill that people typically are not always willing to do the house keeping even though the benefits may be established. Most people want to get away with as much less meta-work as possible. Alan Wexelblat and Pattie Maes (1999), for example, built a navigational system and presented users with several tools to aid collaborative information foraging, but ultimately found that several of their hypotheses proved to be

false as the users didn't bother to use them.

Thomas W. Malone (1983) analyzed how people organize files in offices and found that there are two kinds of people: *Pilers* and *filers*, people who bothered to put in the additional house-keeping with the long term results in mind and people do not bother to do the house-keeping. It is interesting to note the observation of Andy (2009) regarding the bi-modal distribution of types of users: *bingers* and *addicts* when analyzing tab usage data from web-browsers; the habitual tabbers who have more tabs open most of the time and the spiky tabbers who open a large number of tabs but end up closing most of them within a relatively shorter span of time.

However, I think that people have a combination of both tendencies in different ratios. The ratio, I believe, tends to shift for the same person depending upon his energy, mood, etc. So an ideal solution must enable the user to put in just as much effort that he can afford to (or is willing to) at that time and still be useful without deteriorating significantly in usefulness. It should accommodate the variability in user tendencies.

3 Affordances for continued use of information management tools

In particular, I recommend that developers try and rank their tool, with their target audience in mind, along the following various dimensions listed by Jones et al. (2001).

Functional dimensions of information organization tools:

- Context (ability to retain relevant context)

- Reminding (ability to remind the user)

- Portability

- Ease of maintenance

- Ease of integration into existing work-flow

- Communication (ability to share selected information with others)

- Currency (ability to remain relevant and up-to-date)

Bernstein et al. (2008) list a similar set of requirements for an information organization system following a field study:

1. Lightweight capture: lower overhead (*the minimum effort rule*)

2. Visibility and reminding

3. Flexibility in terms of enabling users to devise their own organizational systems

4. Mobility and availability

Jones et al. (2008) identified top 5 reasons that lead to users abandoning information management systems over a 6 month study:

1. Visibility

2. Integration

3. Co-Adoption (affordance for people who work as a group to adopt the system)

4. Scalability (with respect to the type or amount of information that the user wants to process)

5. Return on Investment

The recurring themes identified through the different field studies is interesting to observe. Developers should strive to make their systems extendable as a considerable percentage of users are likely to have their own work-flow that they have grown accustomed to and the work-flow of one user is likely to

differ significantly from another. Hence users should be able to extend the system and adapt it to fit their work-flows, for example through a plug-in system. Users should be able to create their own schemas to fit their data capture and it should be trivial to share schemas with other users who would be interested in reusing them for their own purpose, but do not have the skills to create them.

3.1 Why bother with a Visualization ?

Does the problem justify creation of a visualization based system ? Can't we get away with a system that did not use a visualization to communicate with the user ? We can have an effective system which does not use a visualization. In fact, some of the systems that we will discuss do not use visualizations. But I believe that an information organization system will be highly efficient if it incorporates a visualization for communicating with the user.

The cycle of switching work-spheres, when you want to switch tasks can be divided into following phases:

1. Know where you are right now, the progress of the current work – get a sense of where you are leaving the task at (saving the current *context*)

2. Get an *overview* of what else is out there that needs to be worked on

3. Choose one that you want to work on next, depending on the urgency, and/or your energy, mood etc (*reminding*)

4. Switch to the task, recover *context* (answers to *what*, *who* and more importantly *why* about the task) and start working on it for a while before repeating the cycle again

I believe that this cycle lends itself well to a 'zooming-out, getting overview + context and zooming-in onto the details' visualization system. Considering this with the requirements of lightweight capture and

that of visibility and reminding as listed by Bernstein et al. (2008) and Jones et al. (2001) into account, I think that the user would be benefited by the improved affordances offered by a visualization system rather than a system without a visualization.

Thomas W. Malone (1983) concludes from his study of people's workspaces that reminding is a subtle but important aspect of an information organization system. He characterizes that the system should use different attributes such as size (of icons for example), location and colour to remind users of the priority, urgency etc and should modify these over time (either increase or decrease the strength) to attract/detract the attention of the user. This further strengthens the argument for the construction of a visualization system that could take advantage of specifications.

3.2 How would an ideal visualization solution look like ?

How would an ideal visualization system created for solving the problem look like ? What would its visual characteristics be ? We could start by comparing various representations that have been used in the field of education for the purpose of recollection (as memory aids) and for the purpose of helping students understand a topic (as cognitive aids). This is applicable to our problem in that users during their sense-making tasks are frequently recollecting prior information (memory aid) and trying to understand new information in light of it (cognitive aid).

3.2.1 Comparison of various visualization formats for representing information

MJ Eppler (2006) performs a comparison of different knowledge visualization representations for the purpose of using such representations in education. We can learn a considerable amount from his work for the purpose solving our problem. The visual representations can be divided into two categories based on their emphasis. For examples with representations such as node link diagrams the focus is on

the individual items and their relationships. Whereas in representations like Treemaps, Radar charts, Venn & Euler diagrams, the focus is on the big picture. Hence for our purpose, it is suitable to have a node-link representation as the primary visualization and augment it with a representation that provides the big picture.

Following is a table summarizing the discussion by MJ Eppler (2006). The table has been adapted to illustrate the strengths, weaknesses and subsequently the suitability of the different representations for our problem.

	Emphasis on	Visual representation	Structured	Cost of transferability	Cost of creation, modification	Better suited for
Concept maps	Individual - relationship	Node-link	More	Lower	Higher	Mass Communication
Mind maps	Individual - relationship	Node-link	Less	High	Lower	Personal, individual solutions
Conceptual diagrams	Big picture	Tree-map	More	Lower	Higher	Mass Communication
Visual Metaphors	Big picture	Misc Graphic structures	Low - Medium	Lower [a]	Medium – High [a]	Mass Communication

*Diagram is colour-coded to indicate applicability/suitability of the parameter for our problem.*

Legend: More suited, Less suited, Neutral with respect to our problem under consideration.

Note [a]: Visual metaphor is a powerful but strange beast. It is powerful in the sense that it takes advantage of metaphors, existing internalized concepts & relationships, and utilizes the power of association and similarity to convey information. But care should be exerted in constructing the metaphors as there is the danger of the same metaphor being understood differently and then there is also the danger of not being understood at all if the metaphor relies on specific domain expertise or on knowledge that is not too common.

3.2.2 To mind-map or not ?

The idea of using hierarchical map structures as externalized representations of user's mental models has been explored about a decade back in the social context [Alan Wexelblat and Pattie Maes, 1999]. The authors argue that the hierarchical site structure constructed out of user's actions is a better reflection of the user's mental model illustrating the relationship between the various pieces of information. Kandogan and Shneiderman (1997) observed that restructuring of visited pages by the users and the ability to save and reload the structure might help them make use of the spatial and visual characteristics of the structure to re-find information more efficiently.
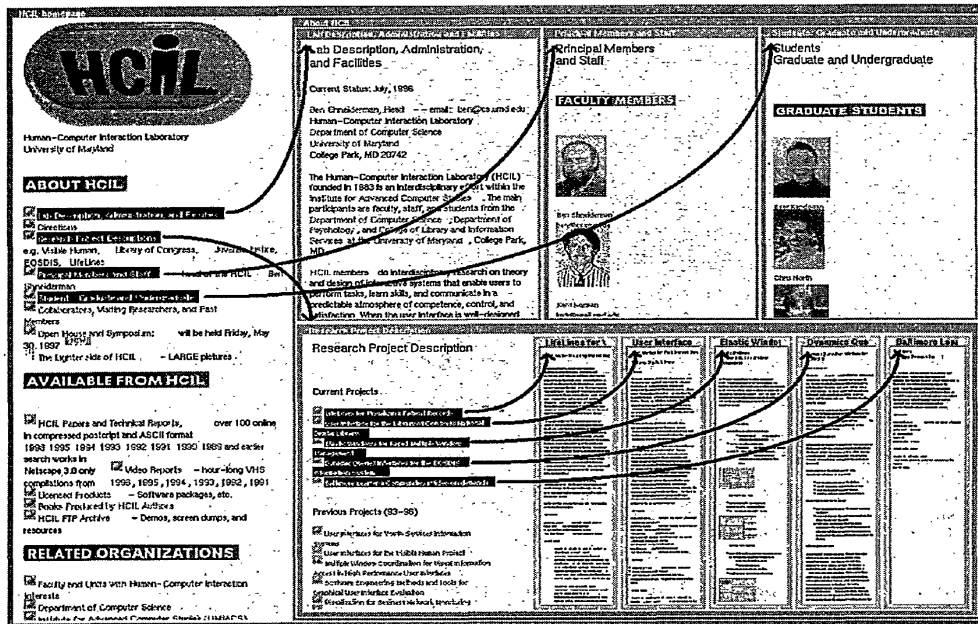
Elastic-windows, a web browser with an implicit hierarchical structure

[Kandogan and Shneiderman, 1997]

Mind-maps, I believe, hence are suited for our purpose as part of the ideal visualization system, due to their low cost of creation and manipulation and given their hierarchical nature which suits the problem domain well to act as a memory and cognitive aid for the user helping them in exploring complex information spaces effectively. We could also consider augmenting the system with other visualizations that are better suited for providing the big picture.

4 Existing systems

We can take a look at the following systems which are trying to solve same or similar problems over the past decade :

- Padprints [Hightower et al., 1998]

- Webforager [Card et al., 1996]

- Footprints [Alan Wexelblat and Pattie Maes, 1999]

- Session-highlights [Jhaveri and Räihä, 2005]

- ScratchPad [Gotz D., 2007]

- InsightFinder [Cheng and Gotz, 2009]

- Presto [Dourish et al., 1999]

- Tasktracer [Dragunov et al., 2005]

4.1 Extension of Text visualization systems ?

We can consider visualization of knowledge ontologies as an extension of text visualization, with more clearly defined relationships between nodes. There is high similarity between the various types of text visualizations and types of knowledge/information visualization tools. For example both tools can be classified as : time-based visualization, attributes-based visualization, and semantic-relationship based visualization. With text visualization systems like Docuburst [Collins et al., 2009] which make use of wordnet, which can be thought of as a semantic ontology, the distinction blurs between both systems.

4.2 Comparison of affordances of existing systems:

Following is a table of what I believe to be reflective of the affordances of the existing systems along various functional dimensions.

	Browser Companion, Plugin or Standalone	Passive/ Active information collection	Lightweight capture	Visibility, reminding	Flexibility	Portability, availability
Footprints	Companion	Passive	High	Moderate	Moderate	High
PadPrints	Companion	Passive	High	High	Moderate	Low
Session Highlights	Companion	Both	High	High	Moderate	Low
Data mountain	Companion	Passive	High	High	Low	Low
Web Forager	Standalone	Active	Moderate	Moderate	Moderate	Low
Presto	Standalone	Active	Moderate	Moderate	High	Moderate
TaskTracer	Standalone	Both	Moderate	Moderate	Low	Low
TimeSpace	Standalone	Passive	High	Low	Moderate	Low
Semantic Treemaps	Companion	Active	Moderate	Moderate	Low	Moderate
InsightFinder	Plugin	Active	Moderate	Low	Low	High
ScratchPad	Standalone	Active	Moderate	High	Moderate	High

4.3 The existing visualization systems can be grouped as:

1. Spatial, depth-cue based

2. Document attribute based

3. Semantic based

4.3.1. Spatial, depth-cue based systems

These systems are based on the metaphor of physical world navigation where we use map, paths, annotations and signposts/landmarks to get around. The systems rely on the spatial memory of humans

and their ability to use depth-cues as a feature to organize documents in three dimensional space.

Examples: Data mountain [Robertson et al., 1998], WebForager [Card et al., 1996].

WebForager and WebBook
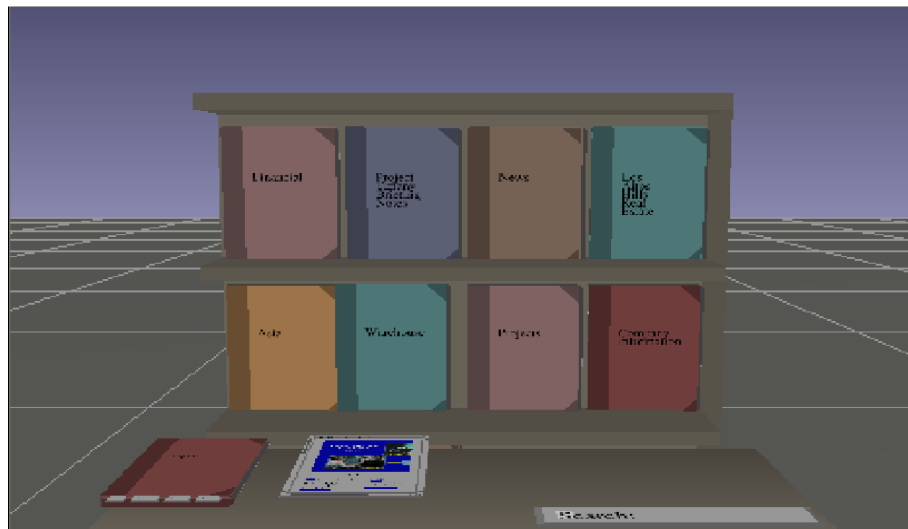
Webforager [Card et al., 1996] for example, helps users organize related webpages into books.

Bookmarks can contain bookmarks inside them to logically delineate several sub-collections.

WebBooks are then placed in shelfs. The users can arrange the WebBooks as they see fit, the spatial

relationship reflecting their mental model.

Data mountain [Robertson et al., 1998] helps users organize their bookmarks in a 3D virtual space.

Users can use cluster similar documents by arranging them closer to one another.

4.3.2 Mozilla Design Challenge

The need for a solution to manage this problem seems to echo through several online communities.

Mozilla, the organization behind Firefox, conducted a design challenge aimed at solving the problem of

information overload in browsers in terms of users struggling with too many tabs, too many bookmarks

and unusable history. Following is the idea sketch from Mozilla's Design Challenge Spring 2009

Proposal winner [Ecaterina Valica, 2009] which illustrates how tabs can be organized and visualized.

The hierarchical tree-structure reflects the interaction history of the user with various pages. The isometric view is similar to the DataMountain in that it allows users to arrange pages spatially in a 3D virtual space.
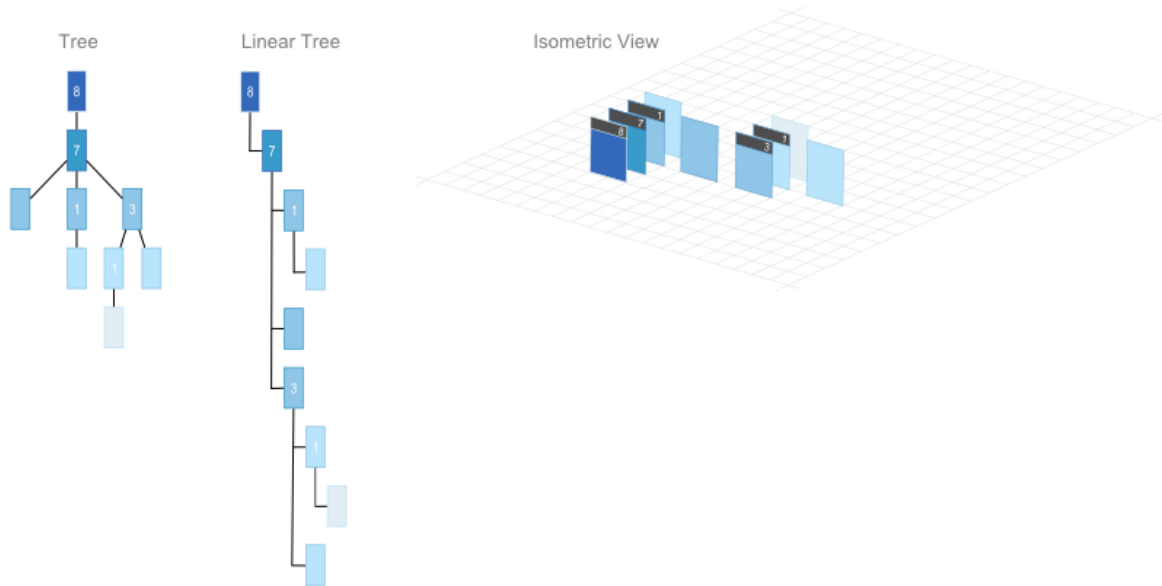

Though these systems might seem straightforward and intuitive on the outset, I have concerns regarding the reliability of our depth-cuing capability and the scalability of our spatial memory. Experiments on utilizing spatial memory to organize information in the virtual world much like as we do in a typical office environment in the real physical world has proved to be less successful [Jones et al., 1986]. Though the informal arrangement of documents in the 3D space might make it easier for the users initially, I doubt whether the system will scale sufficiently with increase in volume/density of information. Limitations on the number of things we can remember spatially (around 10 according to Jones et al. 1986) suggests that spatial memory might be better utilized for storing/retrieving a small number of collections over a shorter span of time. Considering the average number of work-spheres that people typically work on the course of a day, this might work to our advantage if the system is consciously designed with the limitations in mind.

4.1.2. Attribute based systems

These visualization systems take advantage of the various document attributes to organize the

documents.

Presto
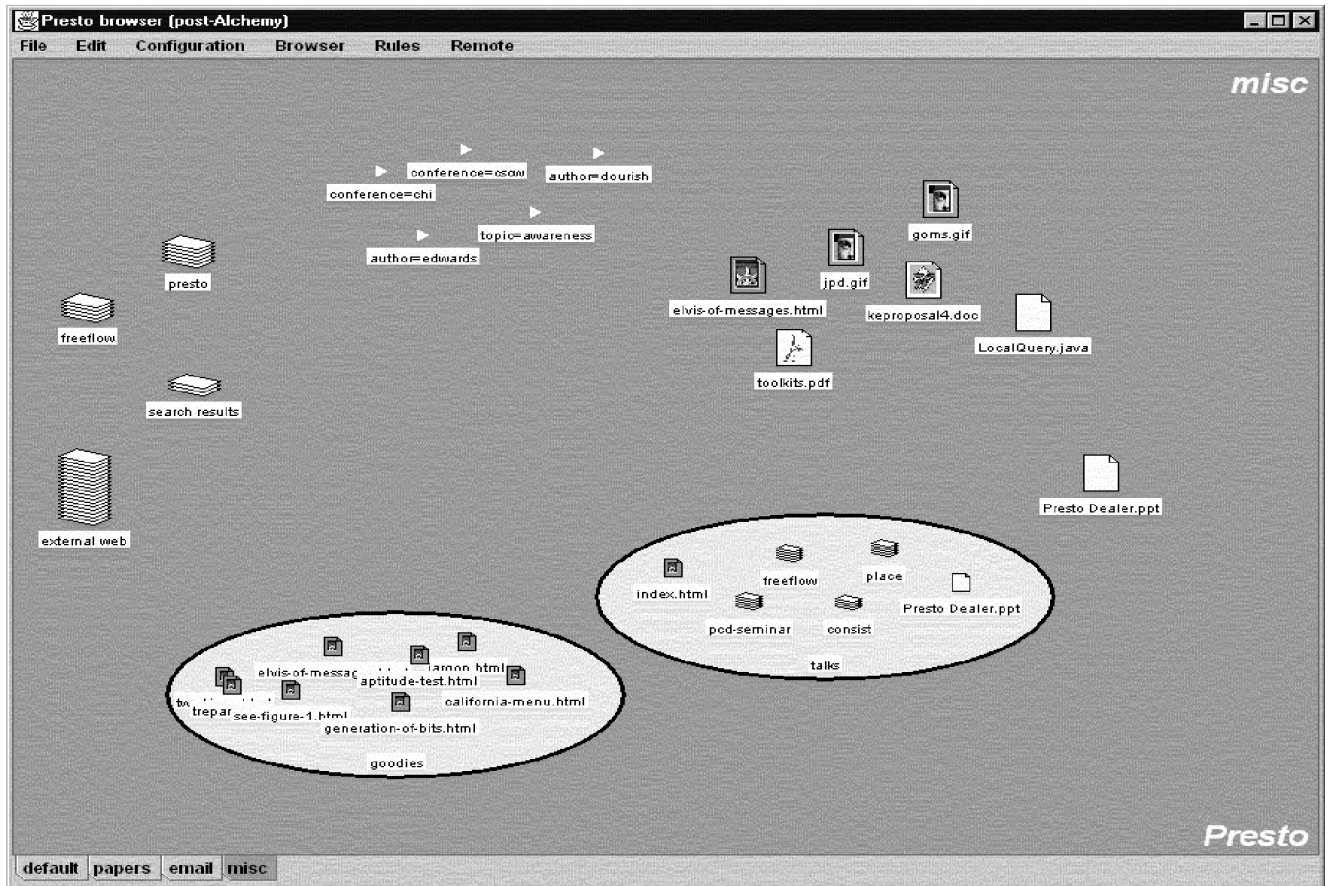
Paul Dourish et al. (1999) Xerox Palo Alto Research Center

Presto is a desktop file organization system created by Dourish et al. (1999). It organizes documents

based on their attributes into workspaces and collections. The piles of documents reflect a collection of

documents that need to be worked on. It follows the physical metaphor of files that need to be worked

on placed on users' desks. When user opens a pile, it expands into an elliptical structure revealing the

documents inside the pile. The attributes which are available for organization are represented by arrows

near the top of the window. Documents which are not grouped into a pile lie scattered around. I think that this system makes good reuse of the physical desk metaphor and take advantages of the file attributes to enable effective organization. But my concerns with this system is regarding its scalability with high volumes of data and the lack of the ability to specify relationships (hierarchical for example) either between workspaces or collections. I think that hierarchical relationships are a good way to manage complexity by organizing related information under sub-structures. It helps us not only to temporarily hide the complexity by not having all information visible, but also to make use of structural relationships that might be better reflection of our mental models, which could enable efficient storage and retrieval of information.

4.3.2 Time based organization

Time is one of the attributes of the information that is used more often to organize, especially by systems that work outside of the browser, since the only information they have access to by default is what resource the user accessed and when.

These systems act either as a proxy for web-browsers (examples: Padprints [Hightower et al., 1998], Footprints [Alan Wexelblat and Pattie Maes, 1999], SessionHighlights [Jhaveri and Räihä, 2005]) or as a client-server based in case of desktop systems, where the server monitors for changes in specific attributes (of files for example) and the client renders them (examples: TimeSpace [Krishnan and Jones, 2005], TaskTracer [Dragunov et al., 2005]).

Padprints [Hightower et al., 1998] for example acts as a proxy server for the web-browser and renders pages visited by the user in a hierarchical structure in a companion window using Pad++. A new node

is added to the structure as when user opens up new pages, provided that it doesn't exist already. User can click on a node in the visualization to bring up the corresponding page. Using Pad++ offers the advantage of scalable zooming in and out to reveal as much context as the user desires and also offers the ability to shrink or collapse subtrees that the user is not interested at the moment.
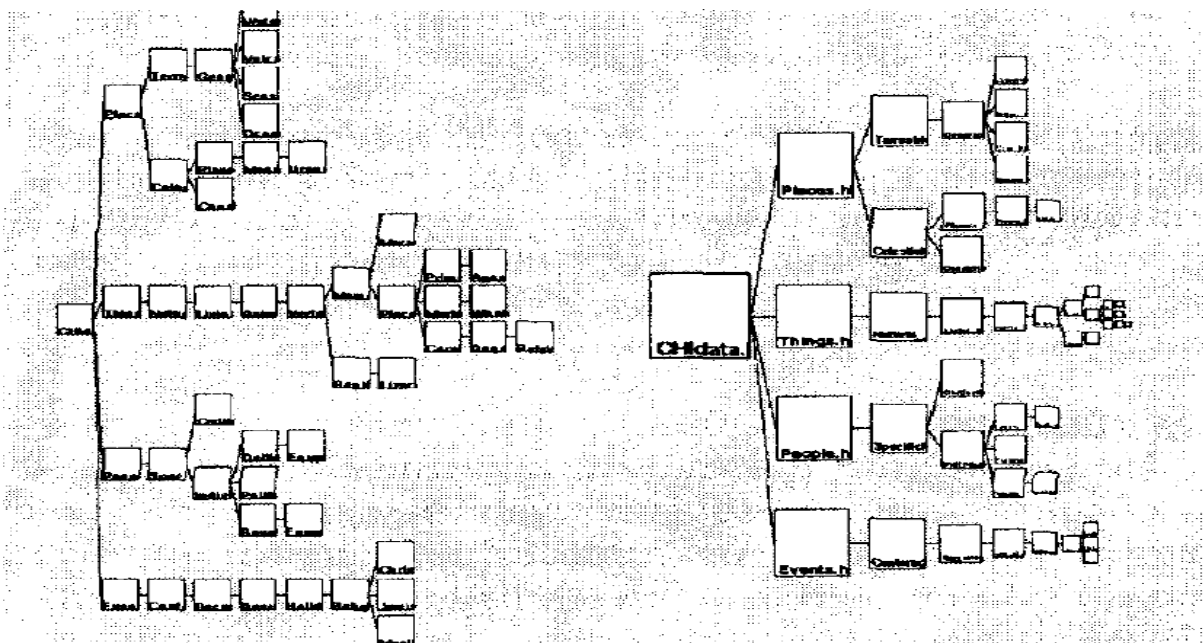
PadPrints (browser companion) visualization

Session Highlights [Jhaveri and Räihä, 2005] is a similar system that renders the thumbnails of the pages visited in the order they visited. Users can save and restore sessions and studies conducted show that users were able to complete tasks spanning multiple-sessions more efficiently using Session Highlights. I think that Session Highlights achieves this by capitalizing on the *Visibility* and *reminding* functionality as identified by Bernstein et al. (2008) by augmenting user interaction with highly visible zoomable icons of pages.
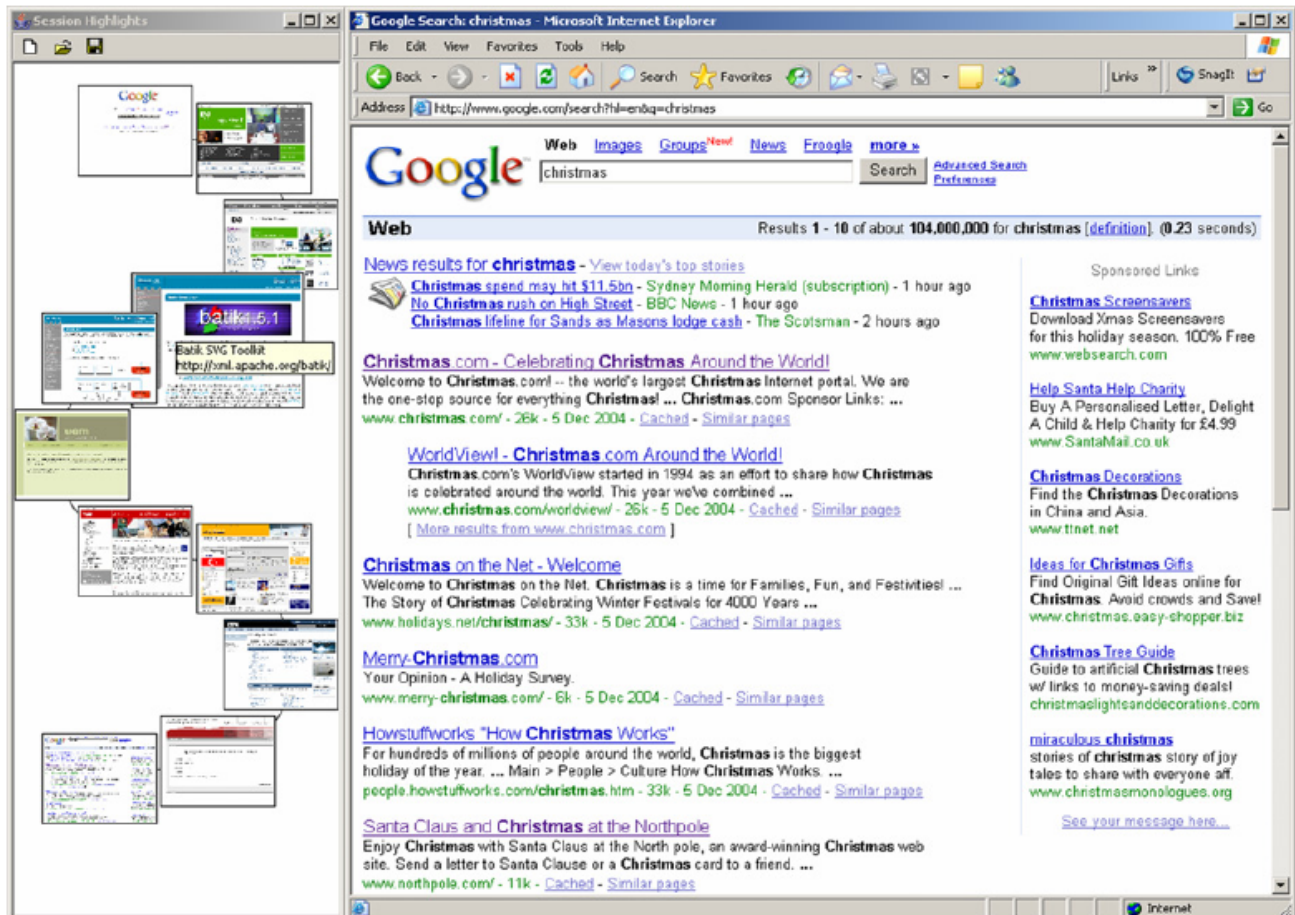
Figure 1: Session Highlights aside a browser window.

Footprints [Alan Wexelblat and Pattie Maes, 1999] is a similar system that acts as a proxy server for the web-browser and renders user's interaction history as a Site-map and a Path-map. Site-map is similar to the hyperbolic tree [Lamping et al., 1995] in that nodes distant from the centre tend to occupy less screen space and this orientation is maintained by dynamically resizing the nodes as user drags nodes away/toward the centre. In order to prevent user from getting lost, a 'Recentre map' function is provided. The Path-map shows the same structure as a hierarchical-list of text entries reflecting the titles of the pages.

4.3.2.1 Social navigation

Alan Wexelblat and Pattie Maes (1999) take the seemingly simple information of user's web history and create a powerful visualization by layering a social context on top of it. The user's own interaction history is shown in the context of the history of other users. Users could glean more information from the paths taken from fellow travellers and this collective foraging, according to the authors, could decrease the cost of information foraging. The resulting personalized site map could also be a better reflection of the user's mental model of the relation between various pieces of information and this according to the authors should be much beneficial to the user than the site-map constructed by the author of the site. The system also adds social information scent to the content of the webpages by adding a number next to the links that reflects the number of people who have followed the link from the page. The system also allows users to comment on both the page and the path itself, which I think is a very good idea. It helps user recover the context easily and answer questions regarding *what*, *why* etc, which was one of the design implication mentioned by Bernstein et al. (2008).
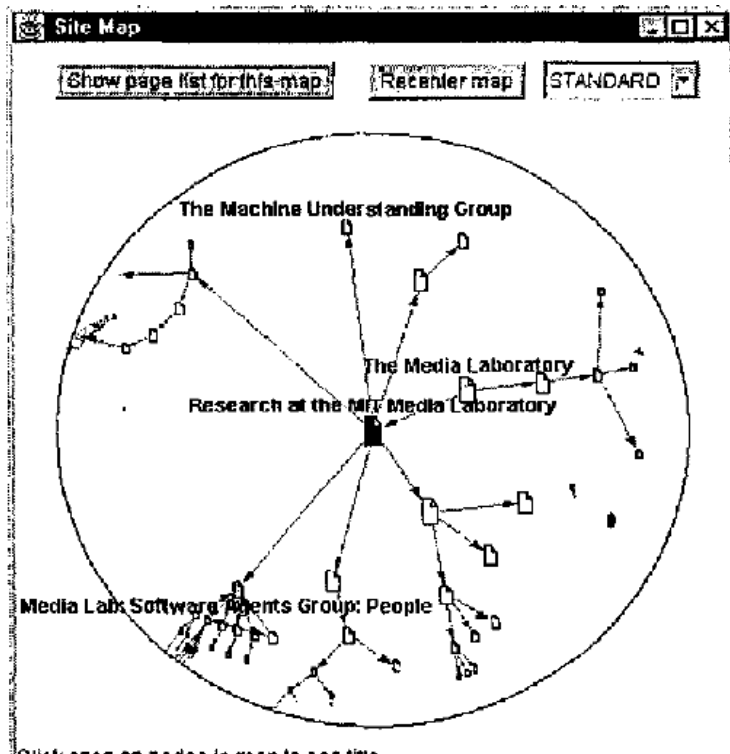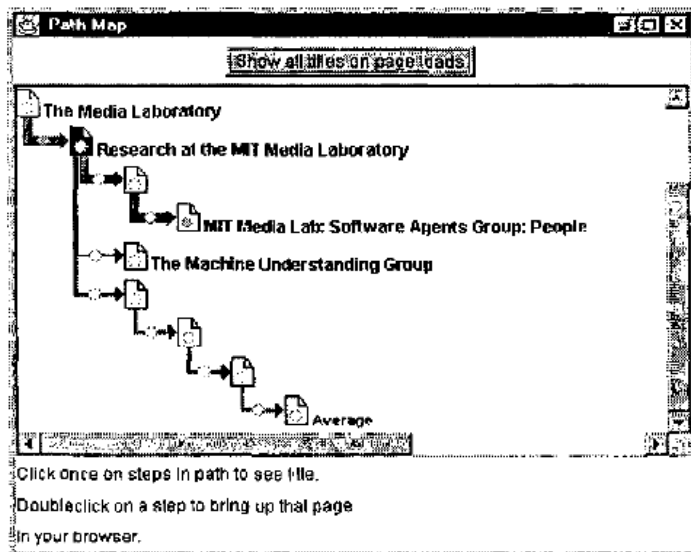
Figure 2 — Footprints Site Map

Figure 3 — Footprints Paths

Footprints (Site map and Path map) used in tandem with a web browser

4.3.2.2 Time flow visualization

TimeSpace [Krishnan and Jones, 2005] is a time flow visualization similar to the Themeriver visualization [Susan et al., 2002] in that it helps user visualize the change of a set of documents over time. Though it doesn't fit well in our problem space, in that it is a desktop application catered towards visualizing changes in files over time, the idea and resulting visualization is interesting enough to warrant a brief discussion.

The central idea being that related tasks and actions performed by the user is clustered in time and helping the user visualize the information could help him in better understanding his information space. I think that TimeSpace could be useful to users in that it helps them to visualize their work-spheres implicitly by laying out the time-based information in contiguous space.
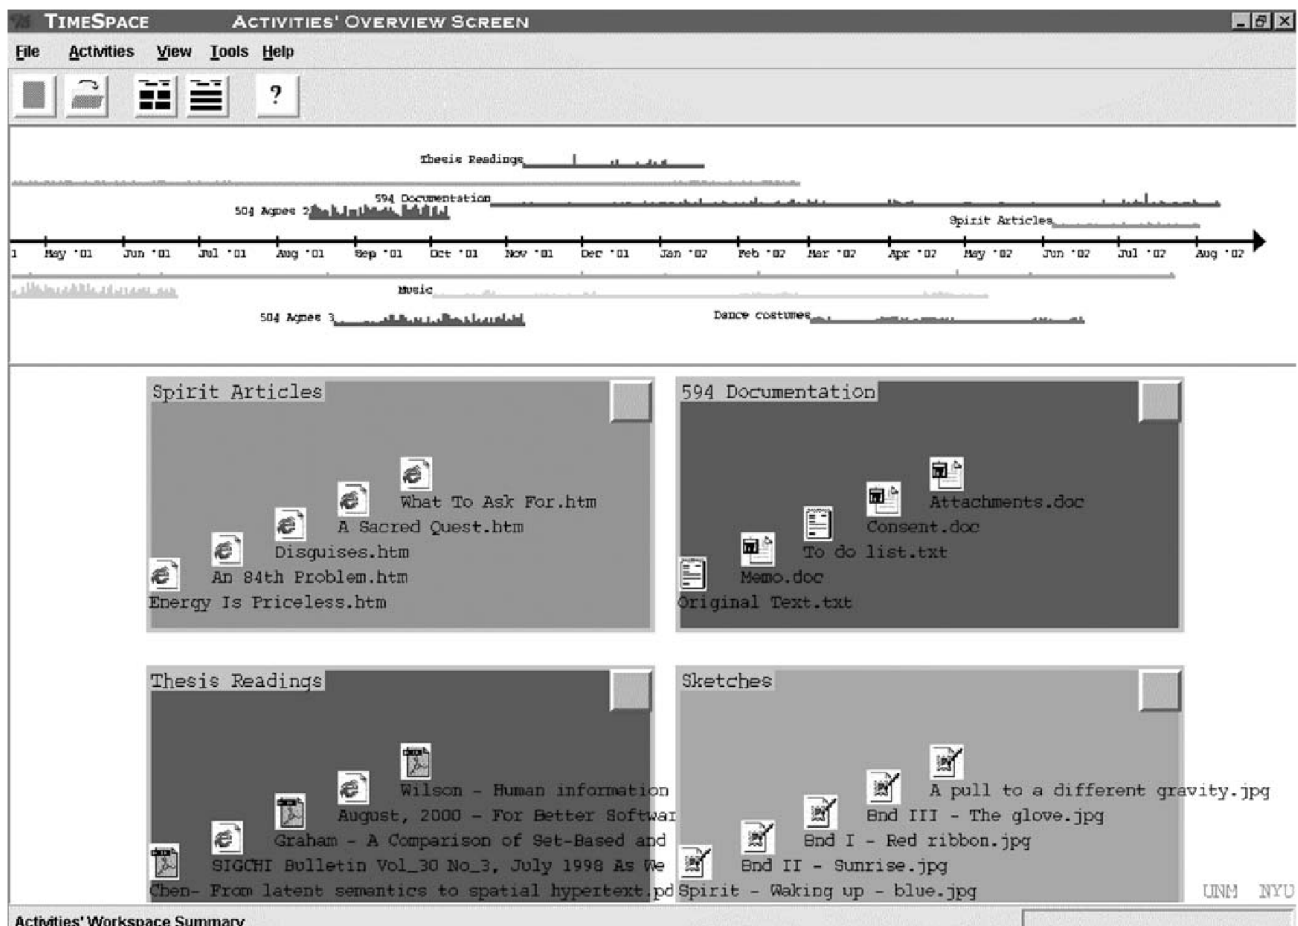
4.3.3 Semantic organization

The ideal solution to this information-drenched condition could be that of ambient findability [Hightower et al., 1998], where finding already found information becomes as trivial as finding new information. One approach is to utilize the semantic relationships between entities to structure them in a way that minimizes the cost of sense-making, storage and subsequent retrieval.

4.3.1 Automatic discovery of contexts and organizing of tasks (Intelligent user interfaces)

InsightFinder [Cheng and Gotz, 2009] is a web browser extension that helps users in their sense-making tasks over the web by highlighting parts of the page that are relevant to their task. It accomplishes this by using a relevancy algorithm which dynamically analyzes the structure and content of the page that the user is viewing with the notes that the user had previously taken as part of his sense-making task.
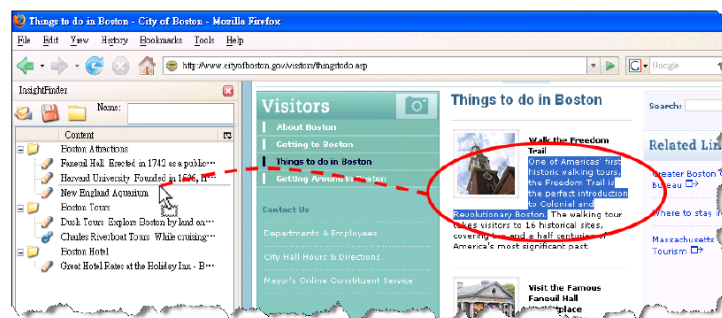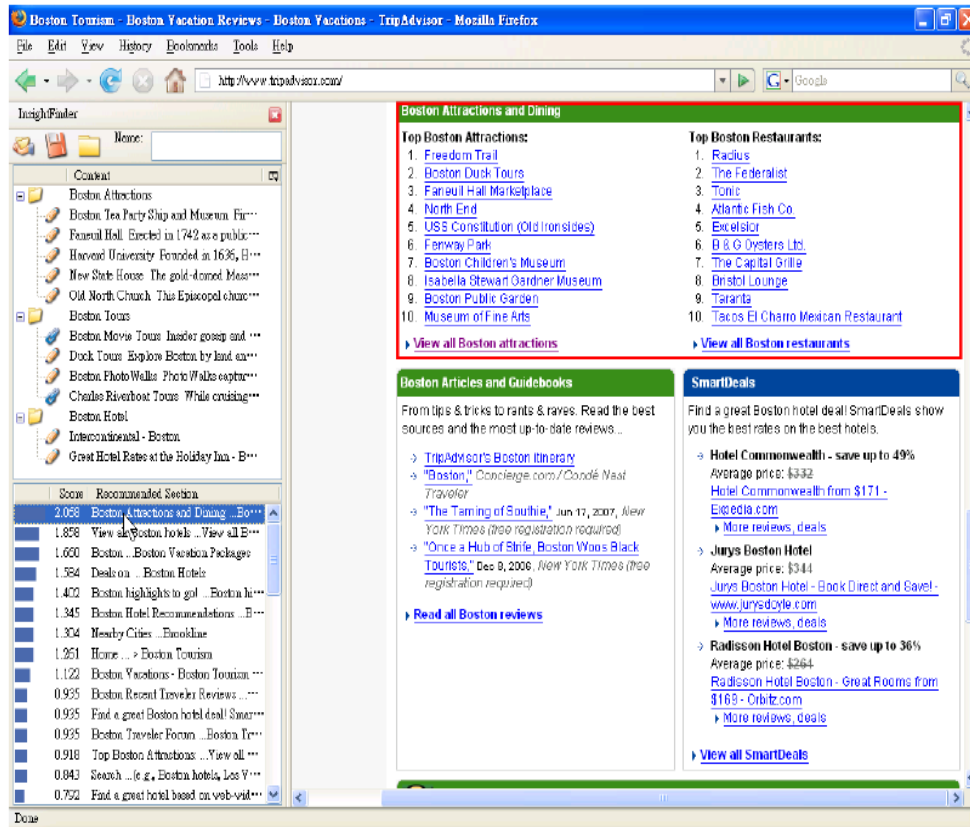

Figure 3. Users can record notes by dragging content fragments (links, images, text, or entire pages) from the browser to folders in the InsightFinder.

ScratchPad [Gotz D., 2007] is a similar tool that offers sense-making support over the web. Users can drag and drop entire web-pages or selected text from a page into the side-bar to create a node. Users can create relations between nodes by joining them and also annotate the relation. This is similar to the feature present in Footprints [Alan Wexelblat and Pattie Maes, 1999] where users could annotate the relationships in the hierarchies. ScratchPad offers sense-making support by ranking the relevancy of existing nodes to the content that the user is viewing currently and then highlighting the nodes accordingly using different shades of green.


User can create a node by dragging and dropping either a complete page or selective text from a page
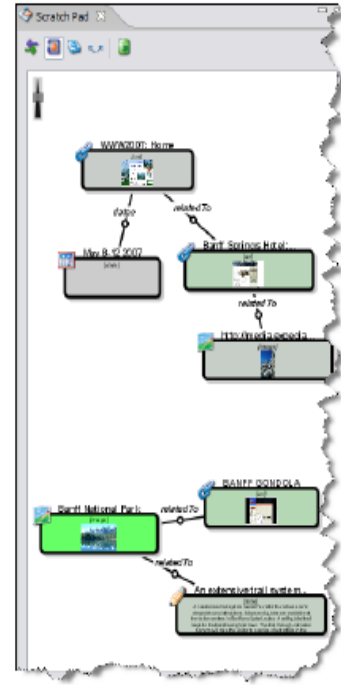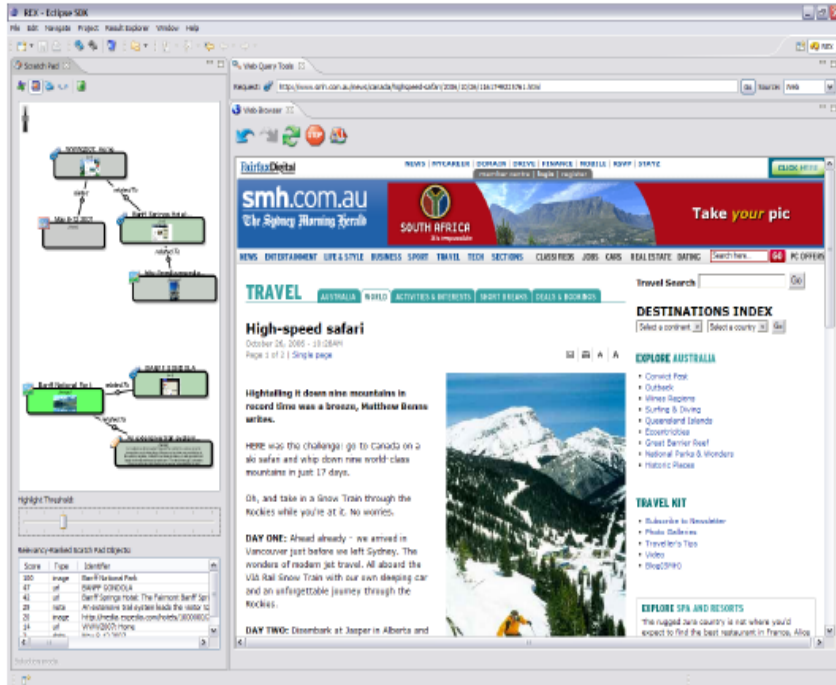
Nodes highlighted in different shades to indicate relevancy of the content of the current page

Dragunov et al., 2005 embed the server part of their application, Tasktracer, in Microsoft applications (Word, Internet Explorer etc) and monitor the applications through clients to discover contexts that the user is working on and take advantage of the information to help user organize his work.
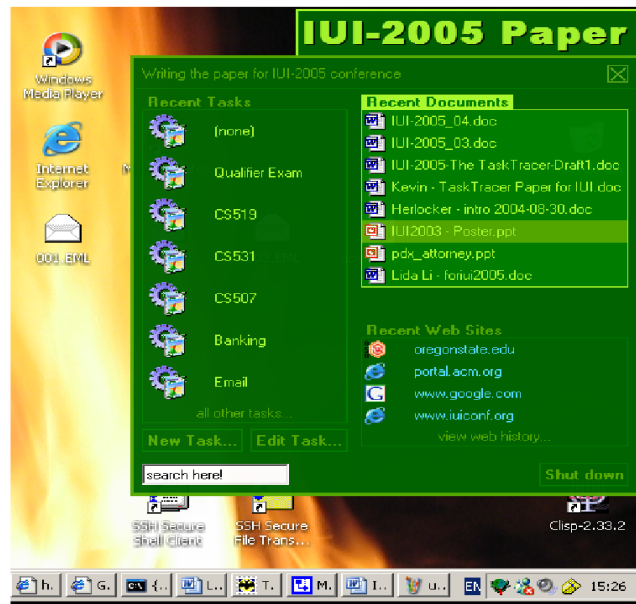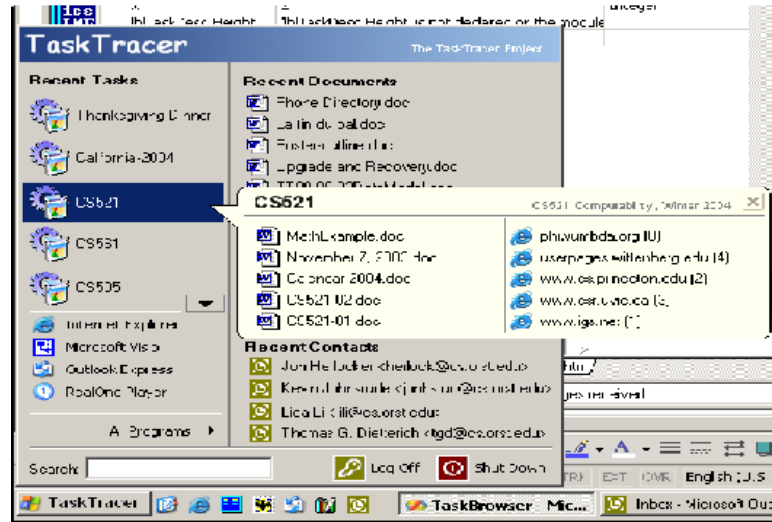
Though the solution seems powerful, given its pervasive nature and integration with and re-use of existing information structures, the biggest problem that I see with Tasktracer is its scalability: that it not portable across computers, platforms and to an extent even across all applications in the same platform that the user might end up using for processing some specific that is part of the information. Relying on a specific technology affects the range of applications that the technique can be applied to and this severely limits the flexibility for the user. This does not go well with *the minimum effort rule*.

Ying Feng and Katy Börner (2002) using Semantic treemaps to categorize and visualize bookmark files. They analyze the semantic relationship between the text of each of the bookmarks and then cluster them based on their similarity using a force directed placement algorithm.
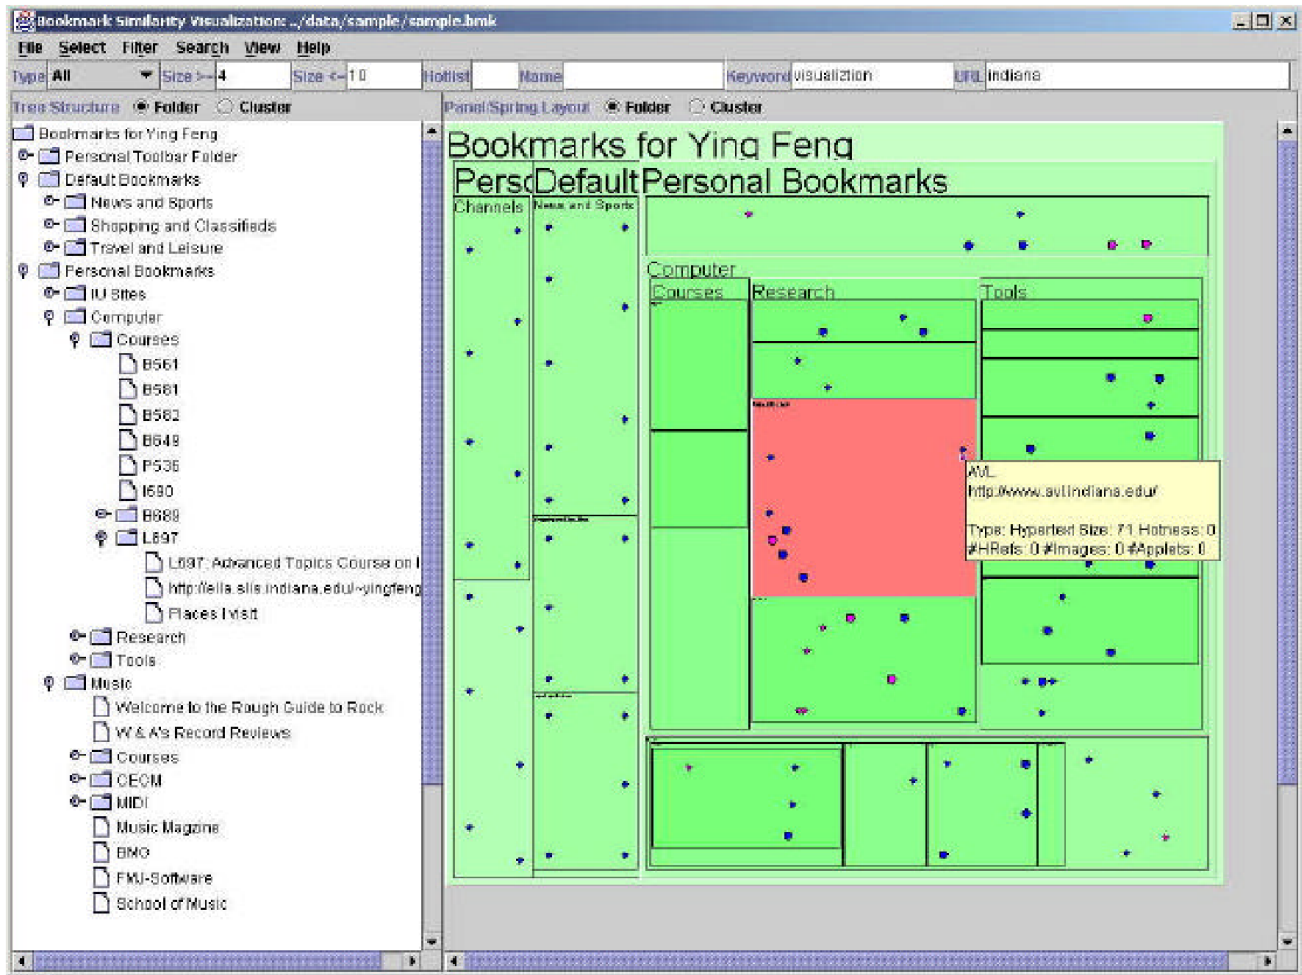
Figure 3. Visualization of a personal bookmark file

They also employ a social layer by constructing a virtual space where users can view their bookmarks in the context of other users, in a semantically clustered layout.
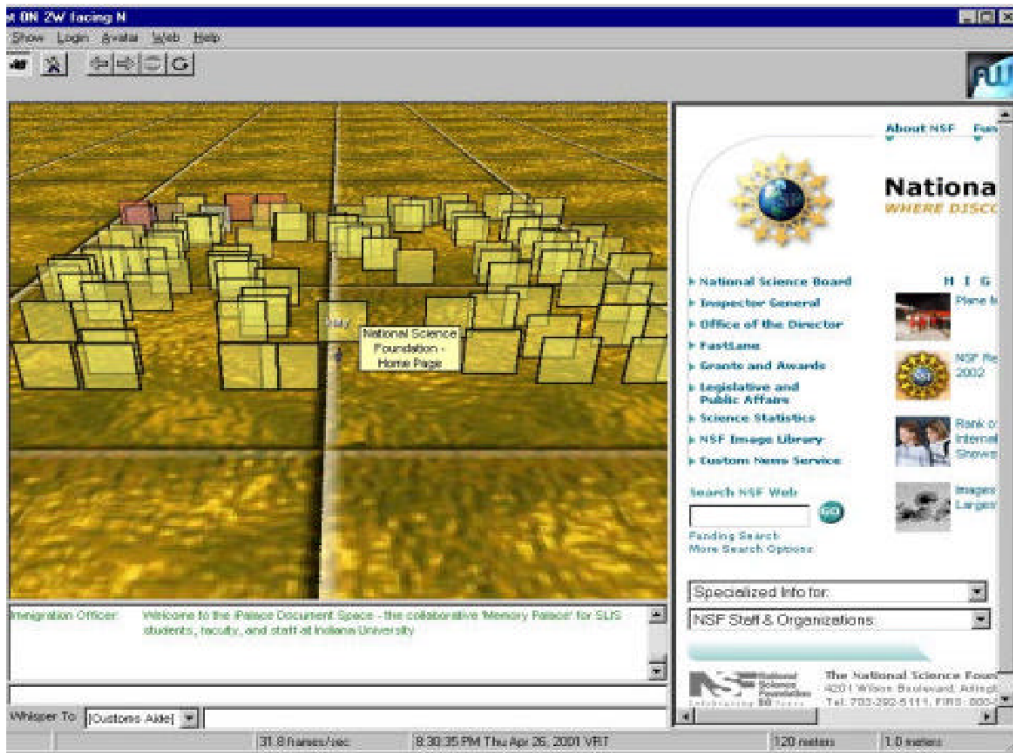
Figure 4. A shared 3-D collaborative bookmark space

4.3.2 Semantic web

If semantic web became a reality and structured semantic information, similar to that of wordnet
[Collins et al., 2009], (wordnet is a hierarchical structure of nodes representing sets of synonyms and
the relation between nodes emphasizing a is-a relationship), building an ideal tool for our problem
space could become a lot easier. Developers could take advantage of the semantic information between
pages to offer powerful passive organizational mechanisms for the users and layering the information
with more useful social ontologies from the social-spheres that the user is interested in.

Magpie [Domingue and Dzbor, 2004] is a semantic browser which utilizes community built ontology to
help the user understand and make sense of the page that is currently being viewed by highlighting
terms corresponding to various ontologies. This could be considered a social extension over context
highlighting systems such as ScratchPad [Gotz D., 2007] and InsightFinder [Cheng and Gotz, 2009],

which help the sense-making task by highlighting content that is more relevant to the information that the user has already amassed as part of the sense-making task. But the usefulness of ScratchPad and InsightFinder is limited to the information that the user has collected. So, ideally the tool would become more accurate as user gathers more information. But I believe that is when the user would need the least help, when compared to the beginning of the sense-making task.

Magpie takes supporting sense-making a step farther, and tries to circumvent the above problem by making use of ontologies constructed by communities that you trust and you belong to. This increases the affordance of the system by making it more flexible and provides higher return over investment.
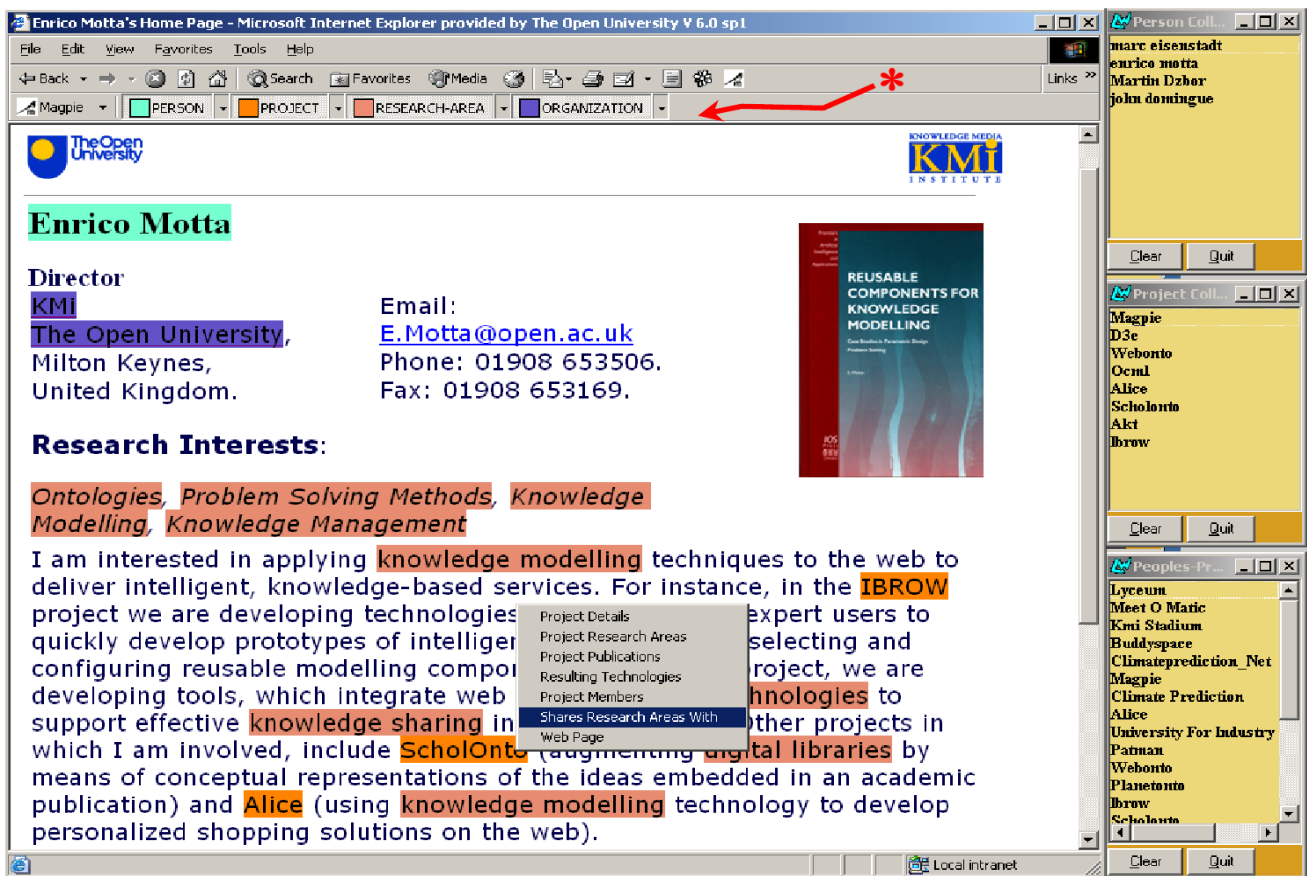
5. Conclusion

Evaluation of information management systems is not straightforward as it might seem to be. Many of the issues crop up only in the long run when the user has had a chance to throw all types of information that he comes in contact with at the system. Lab studies cannot reveal how the system performs with respect to capture and retrieval of information for example when the user is in the *state-of-flow*, for a user can rarely enter such a state in the lab environment. Of the five caveats identified by Jones et al. (2008), I believe that except for *visibility* (which could be evaluated in a relatively short to medium time frame), all other factors could be tested only as part of a long term field study.

1. *Visibility*

2. Integration

3. Co-Adoption (affordance for people who work as a group to adopt the system)

4. Scalability (with respect to the type or amount of information that the user wants to process)

5. Return on Investment

Incidentally, of the systems that we reviewed briefly, those which did a short evaluation reported successful increase in efficiency with the users. For example Session-highlights [Jhaveri and Räihä, 2005] and DataMountain [Robertson et al., 1998] were declared as improving user efficiency after a short routine lab evaluation. But researchers who went to the trouble of conducting a long running field study came up with different results. Users of Footprints [Alan Wexelblat and Pattie Maes, 1999], for example over the course of a year long study were found to be not using the tool in the way that the authors had expected and hence found it more difficult and less efficient than the control system. The authors observed that users resorted to taking the most easy paths rather than the most efficient ones. This, according to me, is one of the most important factors, that developers and researchers of information organization systems should keep in mind when developing or evaluating such a system.

6. Future work

In the lines of Moore's law we could devise an information-glut law, that *every n years, the ratio of information amassed to the information utilized per person will keep doubling.* All our efforts into creating organizational tools for managing the information overload should not just end up allowing users to consume more information than was possible before using the tool. This would result in users falling in the same trap over and over again. As tool authors we should do what we can to avoid this condition.

Some of the things that I would love to see in future tools :

- Increase affordances to help us to stay fit (information-wise) and not accumulate information glut.
- Treat information as an entity and make it more accountable.
- Associate information scent indicating information's nutritional value, similar to the food labels indicating calorie counts (only much more practical and useful):
- Provide preemptive meta-information on the usefulness of a piece of information
  - for example indicate the chances of eventually using a bookmark, while you are bookmarking a resource
    - based on your history with other pieces of similar information
    - based on the information's history with other similar people
- Predict cost of foraging
  - Add information scents, to hyper-links for example, indicating the cost of foraging and consuming the resource
    - based again on the history of both the user and the information matched up with similar entities

- and ultimately prevent us from becoming information gluttons [David Karger, 2009].

References

Alan Wexelblat , Pattie Maes, Footprints: history-rich tools for information foraging, Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit, p.270-277, May 15-20, 1999, Pittsburgh, Pennsylvania, United States

Andy 2009, Getting to Know the TestPilot Tab Usage Data
http://surfmind.com/muzings/?p=505

Bernstein, M., Van Kleek, M., Karger, D., and Schraefel, M. C. 2008. Information scraps: How and why information eludes our personal information management tools. *ACM Trans. Inf. Syst.* 26, 4 (Sep. 2008), 1-46. DOI= http://doi.acm.org/10.1145/1402256.1402263

Card, S. K., Robertson, G. G., and York, W. 1996. The WebBook and the Web Forager: an information workspace for the World-Wide Web. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Common Ground (Vancouver, British Columbia, Canada, April 13 - 18, 1996). M. J. Tauber, Ed. CHI '96. ACM, New York, NY, 111-ff. DOI= http://doi.acm.org/10.1145/238386.238446

Cheng, W. and Gotz, D. 2009. Context-based page unit recommendation for web-based sensemaking tasks. In Proceedings of the 13th international Conference on intelligent User interfaces (Sanibel Island, Florida, USA, February 08 – 11, 2009). IUI '09. ACM, New York, NY, 107-116. DOI= http://doi.acm.org/10.1145/1502650.1502668

Collins, Christopher; Carpendale, Sheelagh; and Penn, Gerald. DocuBurst: Visualizing Document Content using Language Structure. Computer Graphics Forum (Proceedings of Eurographics/IEEE-VGTC Symposium on Visualization (EuroVis '09)), 28(3): pp. 1039-1046, June, 2009

David Karger, Information Glut, or Information Gluttons? MIT CSAIL Research blog. November 2009
http://groups.csail.mit.edu/haystack/blog/2009/11/17/information-glut-or-information-gluttons/

Domingue, J. and Dzbor, M. 2004. Magpie: supporting browsing and navigation on the semantic web. In *Proceedings of the 9th international Conference on intelligent User interfaces* (Funchal, Madeira, Portugal, January 13 - 16, 2004). IUI '04. ACM, New York, NY, 191-197. DOI= http://doi.acm.org/10.1145/964442.964479

Dourish, P., Edwards, W. K., LaMarca, A., and Salisbury, M. 1999. Presto: an experimental architecture for fluid interactive document spaces. ACM Trans. Comput.-Hum. Interact. 6, 2 (Jun. 1999), 133-161. DOI= http://doi.acm.org/10.1145/319091.319099

Dragunov, A. N., Dietterich, T. G., Johnsrude, K., McLaughlin, M., Li, L., and Herlocker, J. L. 2005. TaskTracer: a desktop environment to support multi-tasking knowledge workers. In Proceedings of the 10th international Conference on intelligent User interfaces (San Diego, California, USA, January 10 - 13, 2005). IUI '05. ACM, New York, NY, 75-82. DOI= http://doi.acm.org/10.1145/1040830.1040855

Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., and Robbins, D. C. 2003. Stuff I've seen: a system for personal information retrieval and re-use. In *Proceedings of the 26th Annual international ACM SIGIR Conference on Research and Development in informaion Retrieval* (Toronto, Canada, July 28 - August 01, 2003). SIGIR '03. ACM, New York, NY, 72-79. DOI= http://doi.acm.org/10.1145/860435.860451

Ecaterina Valica, Winner Mozilla Design Challenge, 2009, http://risherry.blogspot.com/2009/02/mozillas-design-challenge-spring-2009.html

Kandogan, E. and Shneiderman, B. 1997. Elastic Windows: evaluation of multi-window operations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, United States, March 22 - 27, 1997). S. Pemberton, Ed. CHI '97. ACM, New York, NY, 250-257. DOI= http://doi.acm.org/10.1145/258549.258720

González, V. M. and Mark, G. 2004. "Constant, constant, multi-tasking craziness": managing multiple working spheres. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems

(Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM, New York, NY, 113-120. DOI= http://doi.acm.org/10.1145/985692.985707

Gotz, D. 2007. The ScratchPad: sensemaking support for the web. In Proceedings of the 16th international Conference on World Wide Web (Banff, Alberta, Canada, May 08 - 12, 2007). WWW '07. ACM, New York, NY, 1329-1330. DOI= http://doi.acm.org/10.1145/1242572.1242834

Hightower, R. R., Ring, L. T., Helfman, J. I., Bederson, B. B., and Hollan, J. D. 1998. Graphical multiscale Web histories: a study of padprints. In Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia : Links, Objects, Time and Space---Structure in Hypermedia Systems: Links, Objects, Time and Space---Structure in Hypermedia Systems (Pittsburgh, Pennsylvania, United States, June 20 - 24, 1998). HYPERTEXT '98. ACM, New York, NY, 58-65. DOI= http://doi.acm.org/10.1145/276627.276634

Jhaveri, N. and Räihä, K. 2005. The advantages of a cross-session web workspace. In CHI '05 Extended Abstracts on Human Factors in Computing Systems (Portland, OR, USA, April 02 - 07, 2005). CHI '05. ACM, New York, NY, 1949-1952. DOI= http://doi.acm.org/10.1145/1056808.1057064

Jones, W. P. and Dumais, S. T. 1986. The spatial metaphor for user interfaces: experimental tests of reference by location versus name. ACM Trans. Inf. Syst. 4, 1 (Jan. 1986), 42-63. DOI= http://doi.acm.org/10.1145/5401.5405

Jones, W., Bruce, H., and Dumais, S. 2001. Keeping found things found on the web. In *Proceedings of the Tenth international Conference on information and Knowledge Management* (Atlanta, Georgia, USA, October 05 - 10, 2001). H. Paques, L. Liu, and D. Grossman, Eds. CIKM '01. ACM, New York, NY, 119-126. DOI= http://doi.acm.org/10.1145/502585.502607

Jones, E., H. Bruce, et al. (2008). I Give Up! Five Factors that Contribute to the Abandonment of Information Management Strategies. 68th Annual Meeting of the American Society for Information Science and Technology (ASIST 2008). Columbus, OH.

Krishnan, A. and Jones, S. 2005. TimeSpace: activity-based temporal visualisation of personal information spaces. *Personal Ubiquitous Comput.* 9, 1 (Jan. 2005), 46-65. DOI= http://dx.doi.org/10.1007/s00779-004-0291-x

Lamping, J., Rao, R., and Pirolli, P. 1995. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, United States, May 07 - 11, 1995). I. R. Katz, R. Mack, L. Marks, M. B. Rosson, and J. Nielsen, Eds. Conference on Human Factors in Computing Systems. ACM Press/Addison-Wesley Publishing Co., New York, NY, 401-408. DOI= http://doi.acm.org/10.1145/223904.223956

Mayer, M., Bederson, B. (August 2001) Browsing Icons: A Task-Based Approach for a Visual Web History HCIL-2001-19, CS-TR-4308, UMIACS-TR-2001-85

MJ Eppler, A comparison between concept maps, mind maps, conceptual diagrams, and visual metaphors as complementary tools for knowledge construction and sharing Information Visualization, 2006, 202-210

Robertson, G., Czerwinski, M., Larson, K., Robbins, D. C., Thiel, D., and van Dantzich, M. 1998. Data mountain: using spatial memory for document management. In Proceedings of the 11th Annual ACM Symposium on User interface Software and Technology (San Francisco, California, United States, November 01 - 04, 1998). UIST '98. ACM, New York, NY, 153-162. DOI= http://doi.acm.org/10.1145/288392.288596

Susan Havre, Elizabeth Hetzler, Paul Whitney, Lucy Nowell, "ThemeRiver: Visualizing Thematic Changes in Large Document Collections," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 9-20, January-March, 2002.

Thomas W. Malone, How do people organize their desks?: Implications for the design of office information systems, ACM Transactions on Information Systems (TOIS), v.1 n.1, p.99-112, Jan. 1983 [doi>10.1145/357423.357430]

Ying Feng and Katy Börner, Using semantic treemaps to categorize and visualize bookmark files. In Proceedings of SPIE -Visualization and Data Analysis, volume 4665, pages 218--227, January 2002