Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication

J A Waterworth
Department of Informatics
Umeå University
S-901 87 UMEÅ
SWEDEN
jwworth@informatik.umu.se

Abstract

Dealing with information objects in virtual three-dimensional space is a promising approach to successfully exploiting the burgeoning availability of networked information. The design issues inherent in using space for information exploration are described, and some solutions suggested. The main focus of the paper is user exploration of the World Wide Web (Web) and how the design of 'personal spaces' could get around some outstanding problems such as providing true multi-threaded navigation support, the back-forward problem, and the need of users to casually organise, reorganise, filter and communicate information. A design for a spatial Web users' environment which addresses many of these issues - StackSpace - is presented. The model is described in some detail at the conceptual level, relating the overall metaphor of personal spaces for exploration to specific topics such as multi-threaded navigation, chronology, currency, customisation and communication. StackSpace is essentially an environment for information explorers rather than providers. To investigate the generality of the organisational model behind StackSpace, it was applied to the problem of information provision. InfraSpace uses the same elements as StackSpace, but attempts to apply them to the design of a spatial version of an existing information service - a public place. Because of the individual, situated needs of information explorers, personal spaces are likely to become an increasingly important element of users' interactive environments.

(To be presented at BCS Conference on 3D on the Internet, Bradford 17-18 April 1996)

1. Introduction and Background

The main guiding concept for this paper is the idea that spatial representation and navigation of information can have significant advantages over other forms, given innate human skills for navigating and handling objects in the real world. People act naturally in spatial worlds, finding their way around by using sophisticated, but largely unconscious, sensori-motor skills. Spatial skills are thus deeply encoded in the human mind and are predominantly perceptual rather than cognitive in nature. But space also has meaning for people, and they find spatial locations intrinsically memorable.

1.1 Structures, Worlds, and Users' Views

Different users, and the same user at different times, will have different requirements for how information is made available to them. They will want it presented in different ways, they will want to organise it in different ways, and they will want to communicate it in different ways.

The scope for such flexibility depends on how the information is structured and how that structure is represented and presented to users. We can distinguish four aspects of this: Structure, Attribute, World Model and User View (see Table 1.)

Structure	**World**	**Attribute**	**User's View**
network	cities	geography	universal
tree	forests	organisation	group
non-tree hiierarchy	boats	technology	frequency

Table 1 - Aspects of Information Organisation and Representation

At the most fundamental level, information is organised according to some Structure. It may be a network of linked nodes, as in hypermedia, or a hierarchy of classes and subclasses, as in a taxonomic biology reference. How that structure is conveyed to the user is referred to as the users' model, the interaction model, the interface model, the interface world, the world model, or simply the World. Often, for graphical user interfaces, the users' model is designed around one or more metaphors such as, for example, a desktop, a library, or a diary.

Individual users may need to select what information is presented to them within the overall framework of the users' model of a given structure. Such customisation is made possible by the provision of User's Views (Waterworth, 1995a). While the world is common to all users, user views cater for the needs of individual users. Views serve as'filters on the world' and are convenient mechanisms for the user performing different tasks, fulfilling different roles, operating with various groups of collaborators, and thus requiring information to be presented in a way that meets his or her particular requirements. One user will often have different requirements from another, and Views allow those requirements to guide the display of information to individuals or groups dynamically (see Waterworth, 1995a). Views also play a role in communication between users and as a reminding mechanism for one user interacting at different times or in support of different tasks. The importance of this for the present work is the way in which manipulable 3D entities can be used to convey, and even create, such customised representations of information and experience.

1.2 Why Spatial Interaction?

1.2.1 Acting Naturally

"We perceive and act in our environment on the basis of concepts of objects, persons, and events constructed in large part out of features and relations which are spatial" (Olsen and Bialystok, 1983).

Historically, developments in HCI design have moved from a focus on the linguistic skills of computer users to their abilities to deal with representations of objects, reflected in the transition from command line interfaces to direct manipulation of graphical representations of computer entities by means of the mouse. Spatial representation of information is a further stage in this tendency to capitalise on innate human capabilities for dealing with space and for precise sensori-motor control. Such control depends on close hand-eye coordination, coordination which can be tapped via a wide range of direct manipulation interfaces.

1.2.2 Behaving Automatically

Sensori-motor skills dependent on hand-eye coordination provide the basis for human skilled performance as exhibited in driving cars, playing the piano, typing and, amongst many other activities, various sports and crafts. These skills are amenable to automatisation, which means that they come to require little or no conscious control (once we have mastered the techniques, we can drive on a routine journey while thinking about something else, for example.). Once automatised, these types of skills can be performed with very low error levels. The necessary knowledge is said to be stored in motor memory, and is accessible for (largely unconscious) use, but not for (conscious) recall. Designing a user interface to tap such skills potentially allows users to interact accurately while focusing on their reason for interacting - rather than on the process of interaction itself.

But the skills that are easily automatised are not easily recalled. There is evidence, for example, that when users are focusing on the subject matter of interest to them, they tend not to remember the routes they took to find the relevant material (Strong, 1993). Like playing a well-practiced piece on the piano, however, the knowledge may be in motor memory even if it cannot be accessed consciously by the user in a direct way. But this would only apply with well-practised routes, which would not be the case for information exploration by definition. With such well-practised skills, it is necessary to go through the sequence of actions in the order in which they are habitually carried out. Once the sequence is started, it will usually be executed flawlessly; it can be demonstrated, but not explained.

Aids to memory, such as signposts, landmarks and other distinctive features, can help remind the user of the route taken or the task carried out. Communication and reminding are closely related, because reminding is essentially asynchronous communication with oneself. It is important that the design of navigational environments takes account of these human characteristics by providing support for starting and communication, including reminding.

Acquiring sensori-motor skills involves knowing how to go about doing something rather than knowing what is needed for the task (see Clancey, 1989). This process is amenable to coaching by an expert, where the focus is on communicating how something should be done by critiquing performances (Schön, 1987). These skills, and interfaces designed to tap them, tend to foster an exploratory approach to content, once the basic style of interaction has been mastered. This makes the approach highly suitable for supporting the process of information exploration. Users should not be thought of as learning where things are; rather, they learn how to behave (almost unconsciously) when looking for things.

1.2.3 The Semantics of Space

It has long been suggested that using the three dimensions of space can enhance the display, navigation and access of information bases (e.g. Fairchild et al., 1988; Robertson et al., 1991, 1993). However, simply using three instead of two dimensions does not guarantee good usability. Norman (1993; p.177-178) suggests that spatial organisation of information only works when i) there is a natural, spatial mapping between the items and the spatial location, ii) desired items can be located in a minimum of attempts, iii) the number of different items at any single location is small enough that they can readily be found, and iv) the amount of work required to try a location, scan its contents, and then try another location is small. It is the first of these, the mapping of non-spatial aspects of information onto 3D coordinates, that is both problematic and (potentially) powerful. The problem is one of attaching meaning in a sensible way to spatial dimensions so that a user, confronted with ‘information objects' arranged in space, can purposively explore these objects, and interpret their shapes and locations easily and naturally, in terms of the concepts in which he is interested.

The systems of spatial organisation that we routinely use in the real, three-dimensional world (book shelves, desktops, wall charts, calendars, filing cabinets) are typified by the design aim of extending two dimensions of space while avoiding the third, depth (Norman, 1993; Chapter Seven). Our offices can be seen as a two-dimensional space that is curved around us. We lay items out on the desktops, on the book shelves and on notice boards so that, as far as possible, nothing is located behind anything else. To extend this physically limited area of display we use devices such as filing cabinets which allow many items to be grouped into several compartments often reflecting subtopics displayed linearly, with different topics arranged in different drawers. We seldom move wholeheartedly beyond two dimensions of organisation if we can possible avoid it. This is because a basic problem with using space in the real world is that when three dimensions are used, some items are often hidden behind or beneath others. When we do appear to use the third dimension, it's usually plain laziness: things near the top of the pile are more recent then those near the bottom. In the virtual world we can, in principle at least, avoid the problems of hidden objects by manipulating the transparency of items within the user’s view.

1.2.4 The Mnemonics of Space

A potential additional advantage of using all the dimensions of space meaningfully is that of memorability. Throughout history, visualisations of three-dimensional space have been used to enhance memory (Yates, 1984). These mnemotechnic systems generally work by associating to-be-remembered items with various locations within an imagined, or more usually real, three-dimensional world such as a set of linked rooms in a building, a complicated theatre set, a well-known city structure, and so on. Spatial organisation is the key to the success of such systems; a large list of items can be recalled easily if each is mentally associated with a particular location in a real or imaginary three-dimensional structure. Recall is achieved by taking an imaginary walk through the memory space. These mental mnemotechnic systems are known as 'artificial memory'. The important thing for these methods to work is that the space is very familiar to the person using the method. Unlike remembering routes in motor memory, this scheme for aiding item recall depends on a sense of the space between items, a sense of how items are located in space relative to each other. This is personal to the user and suggests the idea of personal spaces, in contrast to the public places we happen across when following a sequential path.

It seems plausible that spatial interfaces to information systems can themselves serve as artificial memory for their users. Because such systems are artificial but not merely imaginary, they can capitalise on both imagery and spatial memory in the recollection of the location of items of information. We might thus expect that once users have explored the 3D objects and seen what items are associated with what location in the structure, they would experience little difficulty in recalling where to find a particular item the next time it were needed.

2.0 Exploring Space: Design Issues and Solutions

At the current level of technology, it is difficult to consider both fundamental, generic aspects of spatialization for information exploration, and make specific suggestions to improve the current, temporary situation. For example, the World Wide Web and currently available browsers have severe limitations, such as the underlying approach of discrete pages of information, linked in hypertextual fashion, with little or no context. Working within those limitations on, say, better ways of showing interaction history will produce very different design approaches than if those limitations are ignored (on the assumption that they are temporary) and a new way of exploring a hypothetical, futuristic information network is described. Here we focus on the present and near-future, and the general aspects of information access and provision rather than specific applications.

The main issue in designing spatial interface models is not how to render 3D environments realistically, nor how to support user manipulation of the objects portrayed; it is how to stop users feeling lost. This is the familiar “lost in hyperspace” problem made concrete.

2.1 Exploring and Organising Real Spaces

We can think of users as information explorers, gatherers and organisers. The Web increasingly functions as a virtual world-at-large in which we all perform these roles. We can learn lessons for the design of usable virtual spaces from the ways in which we use and explore physical spaces in the real world.

2.1.1 Users as Travellers

Like all travellers in a strange land we have three main problems: Making sense of what we come across, finding things of interest, and communicating with others and sharing what we have found. As travellers in cyberspace, we can benefit from the experiences of other travellers. We can share commented itineries, travellers’ tales and travelogues. It is actually relatively easy to communicate about spatial exploration, since people are used to memorising and talking about locations, describing routes and so on (Dieberger and Tromp, 1993; Yates, 1984). Such records of journeys would need some editable record of the route taken, and facilities for user annotation.

There are serious problems with evaluating the success (or otherwise) of information explorations, given the absence of objective measures. Analysis of travellers’ tales may provide one route to assessing information exploration environments and tools.

As collectors of interesting finds, travellers need bags, bins or other containers to collect things in during the process of exploration. These can be examined in detail later and passed to others. One model of the steps involved in this process of exploration and collection is known by the acronym GROPE - Gather, Review, Organise, Publicise, and Edit. Firstly, we gather what appears, at a cursory glance, to be relevant to our interests. When we have finished collecting, we stop and review what we have found, going through the things we have collected at our leisure. We then seek to organise the materials in a way that makes sense to us, and to communicate our findings with others (the Publicise phase). Finally, on reflection and perhaps in response to feedback from others, we edit what we have organised, discarding items that no longer seem very relevant or important.

2.1.2 How the World Works and Looks

One important aspect of our experience of terrestial space is the presence of gravity. We make use of this when we pile things up vertically on horizontal surfaces, often reflecting a chronological ordering from bottom to top. A pile metaphor for temporary organisations of data was suggested by Mander et al. (1992).

A further consequence of evolving in a world with gravity is that we tend to be happier exploring horizontal surfaces than vertical, which somewhat offsets our tendency to stack things in piles. We cope better with horizontal layouts than with vertical, as witnessed by the difficulty of training pilots to navigate in the third dimension. Our failure to socialise with others living or working on floors below or above us, as compared with those located nearby on the same level, may be another example (although this may also be due to the typically low connectivity between storeys as compared to within them). The consequence for spatial design is that most navigation should be conducted on the horizontal plane and stacks should be vertical, but not too high. In personal workspaces, such as our offices, we make use of walls and other vertical surfaces to display calendars, posters and other reminders of future events or tasks, such as ‘To Do’ lists. These not only remind ourselves of information, they often communicate that information to others or help us to do so.

Many aspects of real space planning and development are relevant to designing virtual spaces, such as pathways, landmarks and public and private places. For an excellent account of these topics, see Lynch (1960). Habitable towns and cities are not simply planned. Rather, there is a meeting between top-down planning and bottom-up ‘evolution’. Successful paths and landmarks, for example, develop with use. Navigation is also assisted by “writing on the world” (Dieberger and Tromp, 1993) of various kinds, such as signposts showing the direction of places and labels showing current location (the name on the side of a building, for example). Another aspect of reading the spatial world is based on knowing that as things get older they look older. Similarly, as collected information gets older it should look its age, aiding identification and later selection (Hill et al., 1992).

The real world is largely composed of objects, some of them moveable but many of them not, situated at particular locations; and, sensibly enough, we tend to navigate with reference to fixed objects. The skills we bring to bear on navigating in the world and moving objects around (and moving around objects) are sensori-motor rather than consciously cognitive. Similarly, HCI is increasingly a matter of designing artifacts that exercise our sensori-motor (perceptual) skills rather than depending on our problem solving (conscious information processing) abilities (Waterworth, 1995b).

One difficult aspect of built spaces is the inside-outside problem. This refers to the fact that when we are inside a built space we lose perception of the context of that place in relation to other places in the vicinity. From the balcony of a sky-scraper we might get a magnificent panoramic overview of the city. But once we are deep inside such a building, it is very easy to lose our sense of orientation to the world at large. The inside-outside problem also applies to three-dimensional information space. If we construct three-dimensional structures that are explorable from the inside as well as the outside, we are likely to lose the contextual overview once we enter into buildings. This is the case with several systems that use the metaphor of buildings and cities (e.g. Andrews and Pichler, 1994: Dieberger and Tromp, 1993: Musil and Pigel, 1993). The structures serve well as identifiers of categories, and from a high altitude provide a good overview of what is available. But once a user zooms down and actually enters a building, he is likely to become disorientated rapidly.

The easiest way to avoid this problem if to have three-dimensional objects in a three-dimensional space, but only allow objects to be viewed and explored from the outside. This, of course limits the complexity of organisation that is possible. Transparency can also be used to circumvent the inside-outside problem. Consider a mole constructing a complex network of tunnels. Although we are standing above the whole system, we can no more get an overview than can the mole deep within a particular tunnel, because we are outside and cannot see past the surface. But if the rendering is reversed with the earth shown as transparent and the tunnels as black, then we can get a tunnel overview by virtue of being elevated above the network.

2.2 Navigating in Context

To navigate successfully, interface designs must match the characteristics of the information system with the needs of its users. This is largely a matter of providing appropriate contexts to prevent user disorientation.

2.2.1 Supporting Multi-threaded navigation

Many current systems are limited in usefulness by allowing the user to follow only one active navigation thread. Obviously, it is advantageous to have processes that might take some time going on in the background, so that the user can get on with other things. This is the general philosophy behind the various versions of agents that have been implemented, and is compatible with the metaphor of a personal servant who does the tasks we have set while we concentrate on the things for which we are needed. But there are few, if any implementations that actually provide any particular support for multi-threaded navigation even when, as with some recent Web browsers, multi-threaded interaction is provided. To be successful in following more than one navigational thread, users need to be able to maintain an awareness of the context within each thread and to be alerted to the contextual links between threads.

Currently, the navigational model with which most “information explorers” are familiar is not rich enough to support much more than the fairly haphazard collection of interesting URLs. The navigational approach with typical Web browsers is depth-first, with the user following a path consisting of a series of nodes to his current “place” somewhere in the Web. At any particular node, the user has the option to choose a link onwards in his journey, or follow his path back from whence he came. After going back he can choose to go forward again to the next node on the most recent path taken. But he cannot go back to revisit an earlier path, and this is a major weakness with current browsers (the ‘back-forward’ problem). If the browser provides multiple active threads (most don’t), the user may have several active exploration paths in progress. But no relationships between the items in the multiple threads are shown. It is as if each thread were a completely separate entity which may as well be running on a different machine.

Deckscape (Brown and Shillner, 1995) is an innovative attempt to circumvent these limitations. Deckscape is based on a stack metaphor, where Web pages are stacked on top of each other with only the top one visible. When the user follows a link to a Web page, the new page appears on top of the stack. As he follows a single path, the stack gets bigger, with the current page visible on top. The user can revisit the pages in a stack, by leafing through the entire stack, by jumping to the top or bottom of the stack, and by choosing a particular page from a separate list of contents. All new pages are fetched in the background, in a separate thread, so that neither following a slow link nor downloading a large file will freeze the entire application - a common problem when navigating the Web.

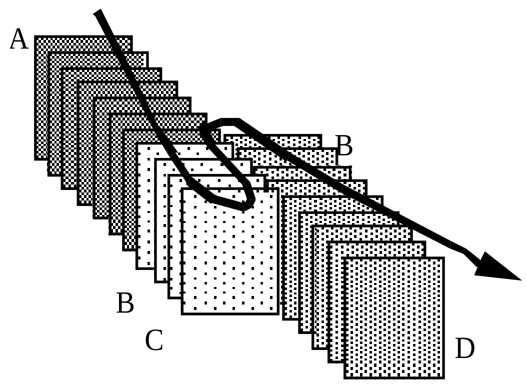

Figure 1 - The back-forward problem

Deckscape keeps all pages until the user chooses to discard them, a facility that can be used to get around the back-forward problem. For example, consider a user who navigates from page A to several other pages including page B and stops at page C (see Figure 1). If he then backtracks to B and moves forward to a previously unvisited page, Deckscape inserts this page (and subsequent ones) after B, and retains the previously-visited pages between B and C. Existing browsers, in contrast, would discard all pages between B and C (and including C itself). Users can also choose to have particular pages displayed separately from the stack and retained there, even if a link on such a page is followed (the path continues on the stack, while the chosen page can be retained for reference whilst exploration continues). And users can form their own stacks comprised of particular collections, such as hotlists.

A weakness with Deckscape is the way the concepts have been visualised. Having only the top-most page on a stack visible means that it looks just like a single page. The separate list of contents does not reinforce the metaphor and selection from it is indirect. The use of real or pseudo-3D to spatialize stacks would add considerable realism, and would tap human spatial abilities.

2.2.2 Chronology, Currency and Focus

Layers, transparency and translucency, and blur are effective mechanisms for providing selective focus. Mechanisms to help users with starting (or restarting) the process of exploration, and to bring new items of information to the notice of the user (monitoring) are also important factors in a successful exploration environment.

If we see information exploration as a journey, we can think of it as divided up into a series of ‘time-slices’ which we can revisit (and do other things with). Stacking is an obvious way of doing this, with the most recent items on the top. Rather than a separate list of what the stack contains, we can label the edges of the slices, if they have the necessary depth. Layers are not limited to serving as time slices, but can be applied wherever selective focus is required. Transparency, translucency and blur have been shown to be very effective in enhancing the perception of layers of graphic display and in isolating objects of interest (Colby and Scholl, 1991). It can also be useful to show history by allowing users to look through from the present into the past, by providing transparent or semi-transparent overlaying of more recent items (e.g. Genau and Kramer, 1995). Lieberman (1994) describes the use of translucency to avoid disorientation when navigating a large display, by superimposing the high level view onto the low level view, thus maintaining context.

Some 3D representations of information have used relative movement to convey a sense of the depth of items in the display, capitalising on our perceptual capacity to use motion parallax as the dominant depth cue (e.g. SemNet; Fairchild et al., 1988). Generally, either the viewpoint of the observer is moved from side to side (or up and down) - the ‘camera position’ - or the display is joggled about in an attempt to create a stronger impression of depth. Moving the camera position is much more successful than moving the display, presumably because this is closer to natural depth perception and because it is hard to concentrate on a jittery display. If head motion detection is available, then depth can be conveyed even more naturally, without having to explicitly manipulate ‘camera’ position. When depth is used in this way, relative movement of layers can also be useful, particularly when there are many time slices to deal with. This allows users to grab a handle attached to a time-slice and move that slice around while the other slices remain fixed. Combined with transparency and viewpoint manipulation, this can be a powerful way of using depth to convey overlaid information (e.g. Silvers, 1995).

Figure 2 gives an impression of the various ways in which time slices might be manipulated to sharpen the focus of user attention.

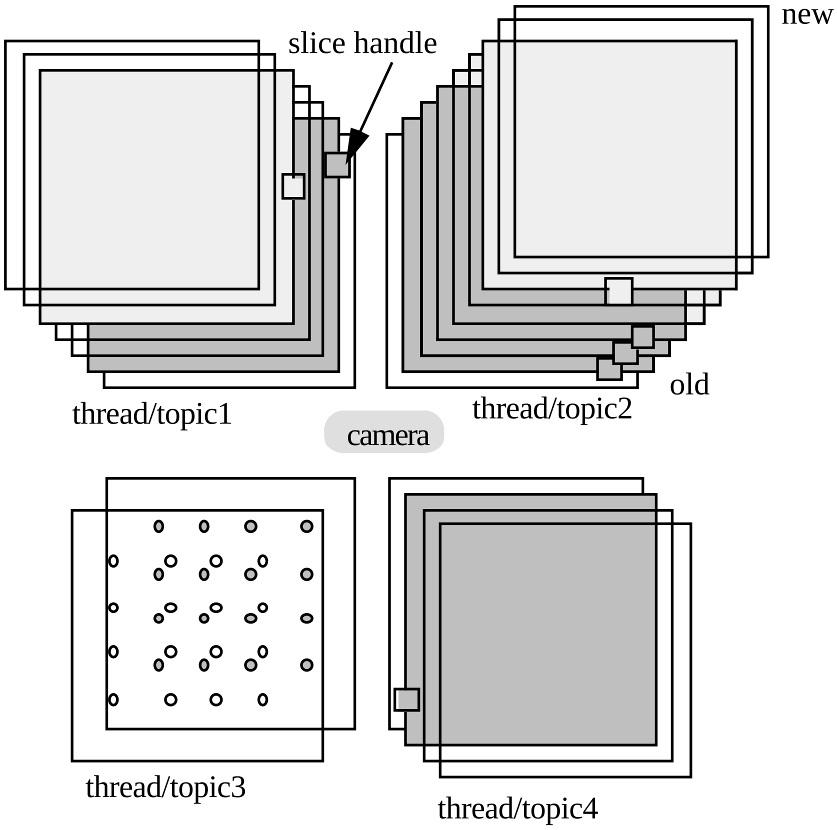

Figure 2 - Transparency, layer handles, and camera position manipulation used to enhance perception of information slices.

The amount of information that is available is expanding at an ever-increasing rate, to the point where we feel overwhelmed by the almost infinite number of nuggets we might discover during exploration. This is the crisis of information provision since, if we define information as that which reduces uncertainty, when we have infinite items of information we effectively have no information. However, if we shift the focus of design from access to information to selection from information, we start to see the user again as a purposive explorer of information space making sense (literally) by the items chosen but also those ignored or rejected. This orientation plays down the importance of exhaustive, precise searching (by machine or human) and stresses satisfying immediate needs and serendipity.

What follows from this is that it is not necessarily desirable to retain all found items. As with human memory, some forgetting by the system might allow us to get on with the current situation with less distraction (Jones, 1989). This idea has been revived by Norman (1993). If we have not used an item of information for some time, we might want it to become less salient. And as time passes and the item remains unaccessed, perhaps it should fade away altogether. Although the prospect of losing information will alarm some users and will not be appropriate for all tasks or situations, temporal pruning will often become essential to avoid information overload.

Related to the issue of old material is that of new. As old information becomes progressively less important, so new information becomes ever more important. Currency is extremely important, and should be marked to alert the user to the fact that he is accessing new information (a process sometimes referred to as monitoring). One popular way of marking new material is with flashing markers. Flashing can be seen as one (rather irritating, but effective) example of an increasing trend to use animation at the interface. More can be done with animation in relation to currency.

3.0 Designing Spatial Worlds for Information Exploration

There are two main aspects of current user interaction where spatial concepts might make a beneficial impact. Firstly, in improving individual users' access to Web-based information in a general way, by designing tools for enhanced information exploration. Such exploration includes finding, selecting from, organising and re-organising information. Secondly, by designing improved interaction models for particular applications to enhance the usefulness of specific information services. These two would comprise personal and public spatial worlds, respectively.

These problems are related, and either one might be expected to benefit the other. Improved user tools would make any particular application easier to use; and a better designed application would be easier, whatever tools were used. So these two approaches might be seen to be compatible and complementary. The process of identifying general problems and possible solutions may also inform that of designing interfaces to specific information services. What follow are design ideas related to these two areas, rather than detailed designs. They can be thought of as illustrative skeletons. Interface ideas are presented but the world is under-developed.

3.1 'StackSpace' - A Personal Space for Exploration of the Web

The motivation behind the design of StackSpace was to provide a personal space where the Web user can follow several exploratory threads (and keep track of them), see the relationships between items in different threads, arrange things according to generic criteria (such as chronology) or personal preference, and communicate not only what has been found but the route by which it was found. StackSpace is an exploration environment and a set of tools for manipulating items of information and the environment itself.

3.1.1 The StackSpace Metaphor

The metaphor on which StackSpace is based - the World model - is that of three-dimensional stacks that contain 'slices' of information. In the context of the Web, slices are in the simplest case Web pages. The principal metaphorical idea is that stacks function rather like piles of paper on a desk, but with extra functionality. Figure 3 shows a 'space' with several stacks of information.

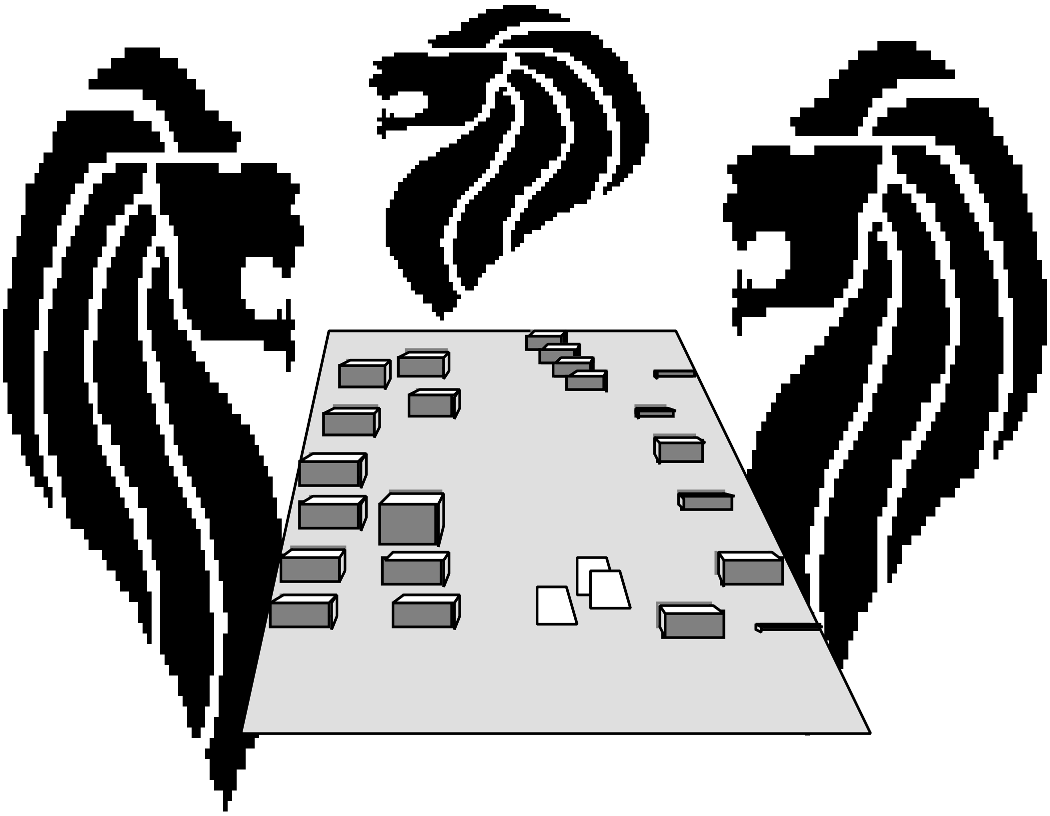

Figure 3 - A Space in StackSpace

The more items a stack contains, the larger it will be. Figure 4 shows two stacks, one composed of six slices, the other of one slice. Items in this world are three-dimensional objects, including the stacks and the slices. Both individual slices and whole stacks can be moved around by the user. The page header is displayed on the edges of a slice to aid recognition of slice contents when they are stacked up. Stacks are automatically built from the bottom up as the user explores pages of information. The side-header scheme for slices means that later slices do not obscure the identity of earlier ones. This also facilitates user-initiated re-organisation and creation of customised stacks.

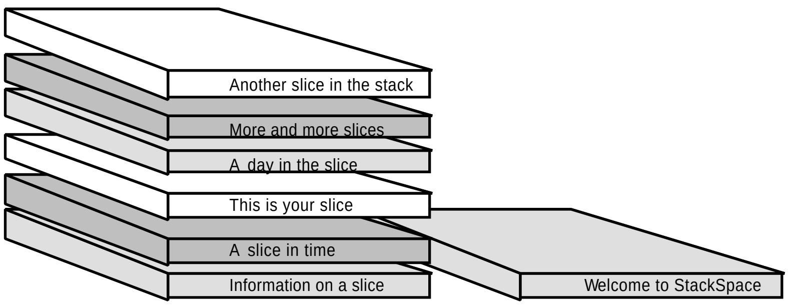

Figure 4 - A six-slice stack and a one-slice stack

Users can see the current slice displayed on the top of a stack. They can revisit any slice in any stack (slices above the chosen one are rendered transparent, or an individual slice can be dragged out of a stack)), after choosing from the titles displayed on the slices. The collection of stacks making up an active interface configuration is termed a Space.

The information structure presented via the StackSpace world is a rather simple one, consisting of quite small, ordered sets of data (pages) arranged in shallow hierarchies. This is appropriate for an environment for gathering and fairly casually re-organising information. The structure is imposed on the underlying unmanageable network that is the Web, and therefore there will be loss of information (though this is not necessarily a bad thing!). The pages comprising the 'leaves' of the hierarchy 'tree' are not really terminals, for example, since they will also point back to other pages at the same and different levels.

This basic structure is expanded to include sets of sets or, in terms of the StackSpace world, stacks of Spaces (where a Space is one or more stacks - see Figure 5). The overall structure of linearised sets that may be arranged in hierarchies is similar to that of other approaches to information organisation, notable NoteCards (Halasz et al., 1987). The users' world, and intended use, is rather different from other hypermedia interfaces based on a similar structure, particularly the present emphasis on dynamic gathering, organising and re-organising, and communication. The assumption about use is that users will be keeping track of several threads represented as stacks, will wish to casually organise materials into a few piles, and will later wish to organise these customised stacks further as spaces that can themselves be stored in a few stacks-of-spaces.

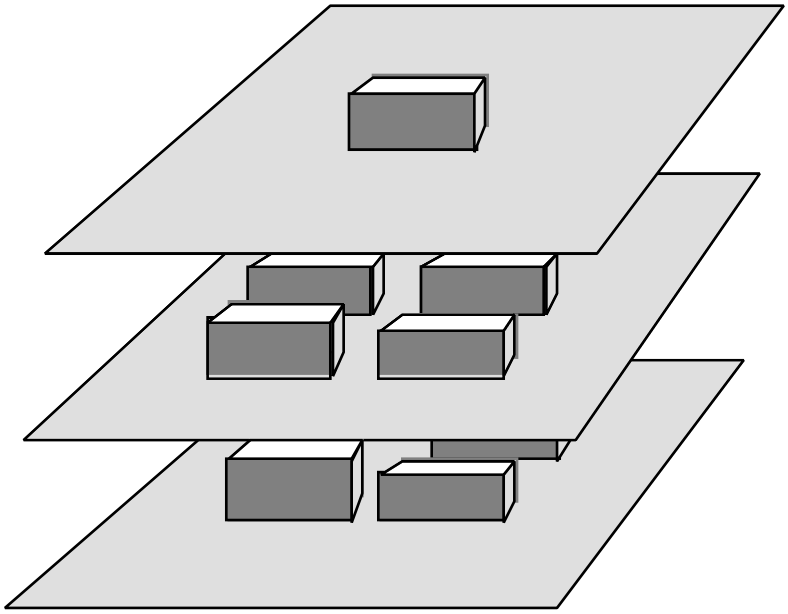

Figure 5 - Spaces as slices, including a single stack space

3.1.2 Multi-threaded Navigation, Chronology and Currency

StackSpace supports multi-threaded navigation with the notion of parallel development of multiple stacks. Instead of opening several windows, as is conventionally done now to follow more than one thread, users can work on more than one stack at once. As they follow links in each thread, a stack develops reflecting the history of that thread. The structure of an automatically-created stack reflects chronology in the same way that piles of papers in our offices do; the further down the stack an item occurs, the older it is in the interaction history. As already noted, there are similarities with DeckScape (Brown and Shillner, 1995), although Deckscape lacks the spatial elements of the current approach.

The limitations on moving back and forward again (back-forward problem) are avoided by allowing users to choose whether to insert slices in the current stack, or preserve them as a new stack, rather than deleting previous "forwards" after a subsequent "back". Figure 6 (left side) illustrates the situation where a user has followed a trail forward from A to C, backtracked to B and then moved forward to D. Existing browsers would now delete the branch from B to C so that the user cannot revisit them easily. StackSpace (right side of figure), in contrast, starts a new stack when the user backtracks to B and then moves on to

D. He can now choose to delete B-C (as existing browsers would, leaving a single stack that does not capture history accurately), have it inserted in the original stack (B-D stacked on top of C), or leave A-C as it is and keep the new stack B-D. The common slice B is linked (see figure) to convey what has happened and allow the user to take whatever action suits him best.

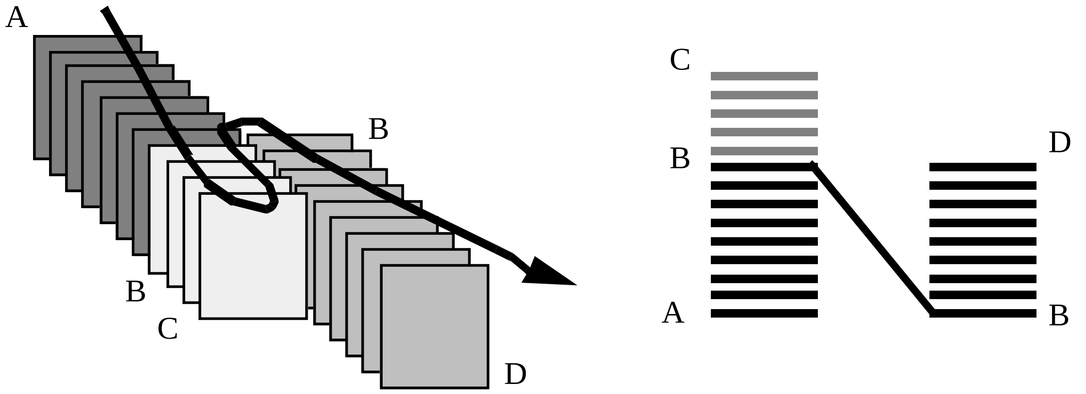

Figure 6 - Dealing with the backwards-forwards problem

The slice on the top of a stack is the current one. When multiple threads are followed, several slices - the top item of several stacks - will be current. These current stacks will by default be arranged across the front of the 3D world, the bottom of the window - i.e. nearest to the user in a world where perspective is used to convey depth.

Stacks that are further away from the user are generally older, and stacks in the distance are older still. But distance from the front is more a measure of relevance to the user than of age. In the flat world of StackSpace, a rather literal cut-off point can be set at a certain level of interest. Stacks that approach this limit will be seen as receding into the background, eventually disappearing over the cliff edge of this flat world to trouble the user no further. The retentive user can extend this horizon to infinity, given enough storage space. The age of items should also be indicated by direct cues; i.e. older stacks and slices should look older than more recent ones. This can be done in various ways, such as progressively deforming the straight edges as the item ages. If required, at a certain age slices and stacks can simply disappear. This is one of the ways of reducing information overload.

When there are contextual relationships between items in different stacks - such as the same slice turning up - these can be shown as ‘bridges’ between stacks, as shown diagrammatically in Figure 7. Links between items on slices can be shown as cords. The situation depicted in Figure 7 might be that two current stacks (in the foreground) A and B have both reached the same Web page (top slice of both stacks) as shown by the bridge between them. The connection from the top of B to part way down C indicates that this same page was also encountered during an earlier session involving stack C. A similar link between two pages of C and D is shown by the bridge between them.

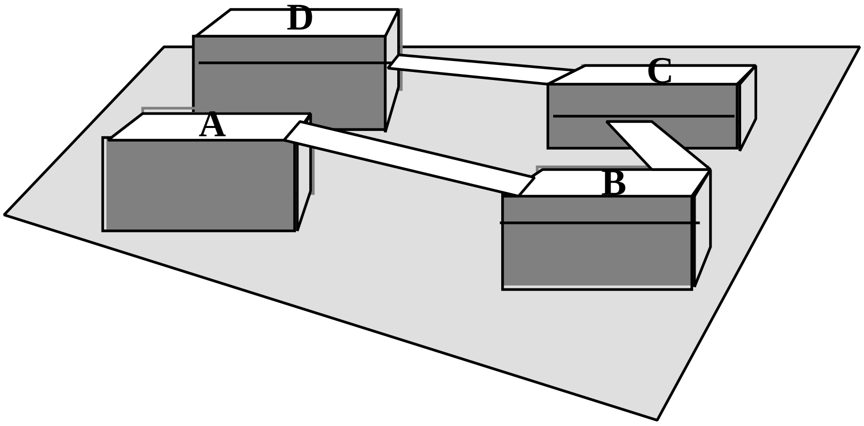

Figure 7 - 3D StackSpace with contextual 'bridges'

3.1.3 Exploration and Agency, Customisation and Communication

Stacks are constructed as users navigate, and as they finish with some stacks and move onto fresh ones, the older ones are replaced in the foreground. Users are also able to launch computer-mediated searches (see Waterworth and Chignell, 1991), and the results of those searches are themselves shown as stacks. In other words, agents can construct stacks, and auto-constructed stacks are one kind of static User View. Users can also manipulate the objects in StackSpace - stacks and slices - to suit their own purposes.

Users can construct stacks by copying or moving slices from other stacks. They can also save stacks, or collections of stacks comprising a Space, as slices in a new stack-of-stacks. A particular configuration of information in one or more stacks (a Space) can function as a static user View. Different stack arrangements can thus comprise different views of the same information space, views that can be communicated with others. As a user navigates, he can select items to be copied to a hotlist stack, a special kind of static view. Since an arrangement of stacks as a space, developed with use, can capture a particular session of information exploration (complete or edited), such spaces can serve as ‘traveller’s tales’ that may be exchanged with others. Again, these are a form of static view. As communication and reminding are closely related (since reminding is asynchronous communication with oneself), traveller’s tales and other static views can also serve as reminders of work in progress or of the context of a particular exploration session.

These kinds of static user view do not fully capture the‘filter on the world’ effect of dynamic views as described in Section 1.1. A dynamic view here would be, for example, one that rendered all slices that did not contain a particular topic, or medium (such as video), transparent. Whatever and whenever stacks were examined this would apply, whereas static views are‘ossified’ - they record how someone organised and selected from a particular set of information at a particular time.

An important aspect of customisation is facilities for marking slices, stacks and spaces. Slices are automatically labelled with the page title. User’s will also want to label stacks, whether they have been created automatically by the user browsing, by an agent or other computer-mediated search for information, or constructed by the user himself. There are many ways in which this could be done. Spaces, stacks and slices should be labelled so as to be identifiable from a distance. It should also be possible to mark slices and stacks temporarily to indicate high relevance. This could be done in several ways, by highlighting the slice or stack (flashing, or especially bright) or by having markers that can be attached to slices and stacks (spinning disc, glowing globe, etc.).

In the 3D world, a viewpoint is, literally, the point in the virtual space from which views are taken, and a perspective is a particular view from a particular viewpoint - a particular filtered view of the world taken from a particular angle. In the simplest case, perspective corresponds to camera angle, viewpoint to camera location. Users can manipulate the viewpoint and perspective to control which stacks are in clear focus. The viewpoint can be moved in six directions: north, south, east and west, and up and down. Figure 8 shows viewpoint movement (combined with changed perspective) and its effect on the StackSpace display.

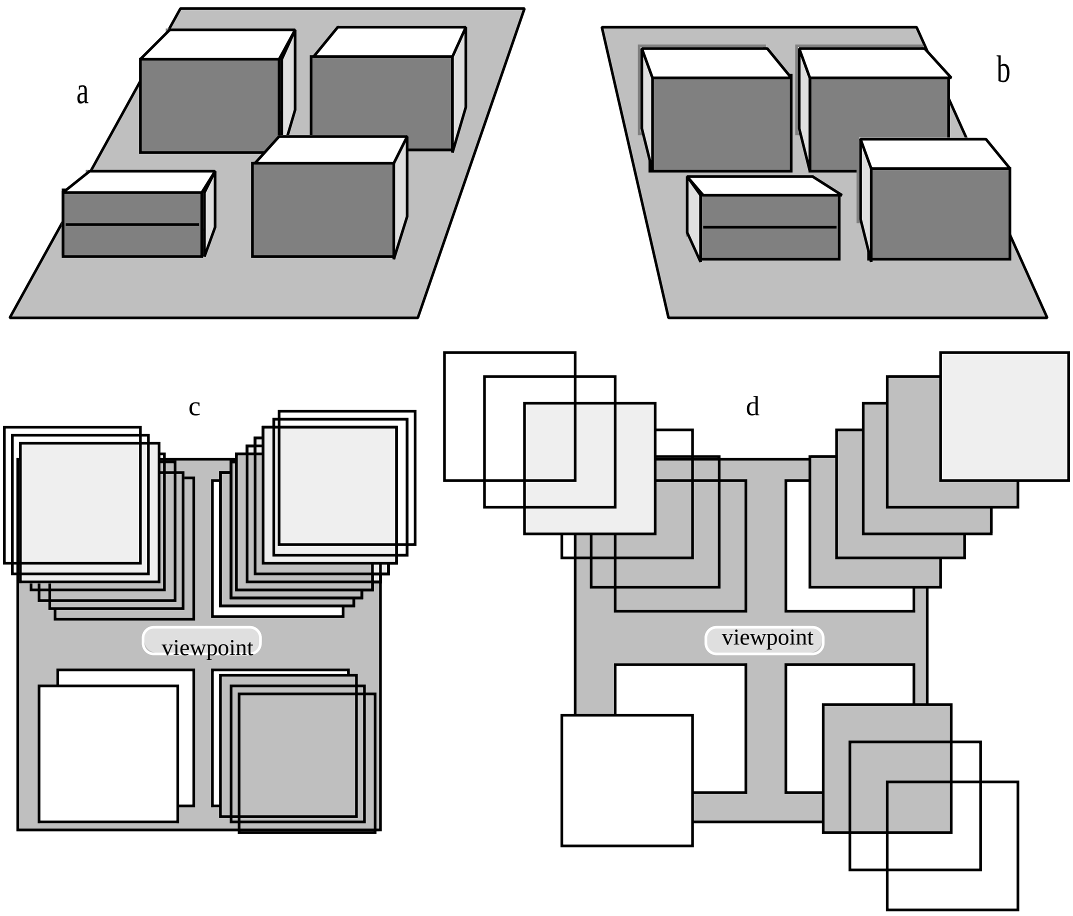

Figure 8 - Moving Viewpoint and Changing Perspective

4. ‘Spatializing’ Existing Applications

StackSpace addressed the needs of the generic‘information explorer’ and was an illustration of how 3D could be used to enhance user interactions with networked information sources such as the Web. In this section, the question of how well such an approach would meet the needs of information provision is considered, by examining a Web-based information service available only to employees of a particular company. Such “intra-nets” are becoming increasingly widespread as primary sources of in-house information.

4.1 InfraNet - An In-house Information Service

InfraNet is an existing Web service providing information for employees of a large company. It is not accessible by the general public. InfraNet has broad scope and yet has a clear purpose and user population, making it a good candidate for re-design on spatial lines.

InfraNet is organised into 12 topic categories set out vertically at the home page:

- World Wide Web
- Security
- Safety
- Company Information
- Company Directories
- Instructions and Forms
- Programmes and Units
- Information Technology
- Information Services
- Personnel
- Newsletters
- Other Information

In terms of the structure of information provided, InfraNet can be characterised as a service that uses twelve top level topics, broken down into a shallow hierarchy of options (a few options - at most five - per level), terminating in linear sequences of pages. This makes it quite suitable, in terms of structure, for realisation as a world of stacked pages and spaces. But it is no accident that in both InfraNet and StackSpace the number of levels and items per level is relative low. This is what makes InfraNet a relatively easily comprehended service, and StackSpace suitable for fairly casual organisation of incoming information according to individual needs.

A major difference between the two is that Stackspace supports exploration of whatever information the user comes across, whereas InfraNet is a designed set of information with a specific purpose. Staff use InfraNet to find information about the company that will be useful to them as employees. Since InfraNet is a widely used service and its content is maintained, the focus of use is very different from that of StackSpace, and therefore there will be significant differences in design. It is an interesting exercise to apply a StackSpace-like solution to this kind of problem.

4.2 'InfraSpace' - Spatialized Design of an Existing Information Service

It was decided that most of the InfraNet categories of information should be retained, and design efforts focus purely on suggesting a spatial interface to the existing information service. The main reason was the heavy use of the existing service by a large number of company employees.

Whatever structure for organising information were chosen, we need 3D entities to represent these major categories. These major objects in the information world of InfraNet should have a recognisable identity when viewed from a distance (for an overview), and a navigable and understandable structure when viewed from closer in. The major categories are represented as Spaces. These objects should be distinctive from each other by graphic design and colour and, to increase distinctiveness, the top-level view portrays them as distinctive 3-D shapes. The hierarchy of world entities in this model is Spaces composed of Stacks, Stacks composed of slices (which may themselves be Spaces) and, at the lowest level, slices that represent individual Web pages (see Figure 9). In InfraSpace, spaces only exist at the level of the major categories of information. Objects in this world have no enclosed interior, thus avoiding the inside-outside problem.

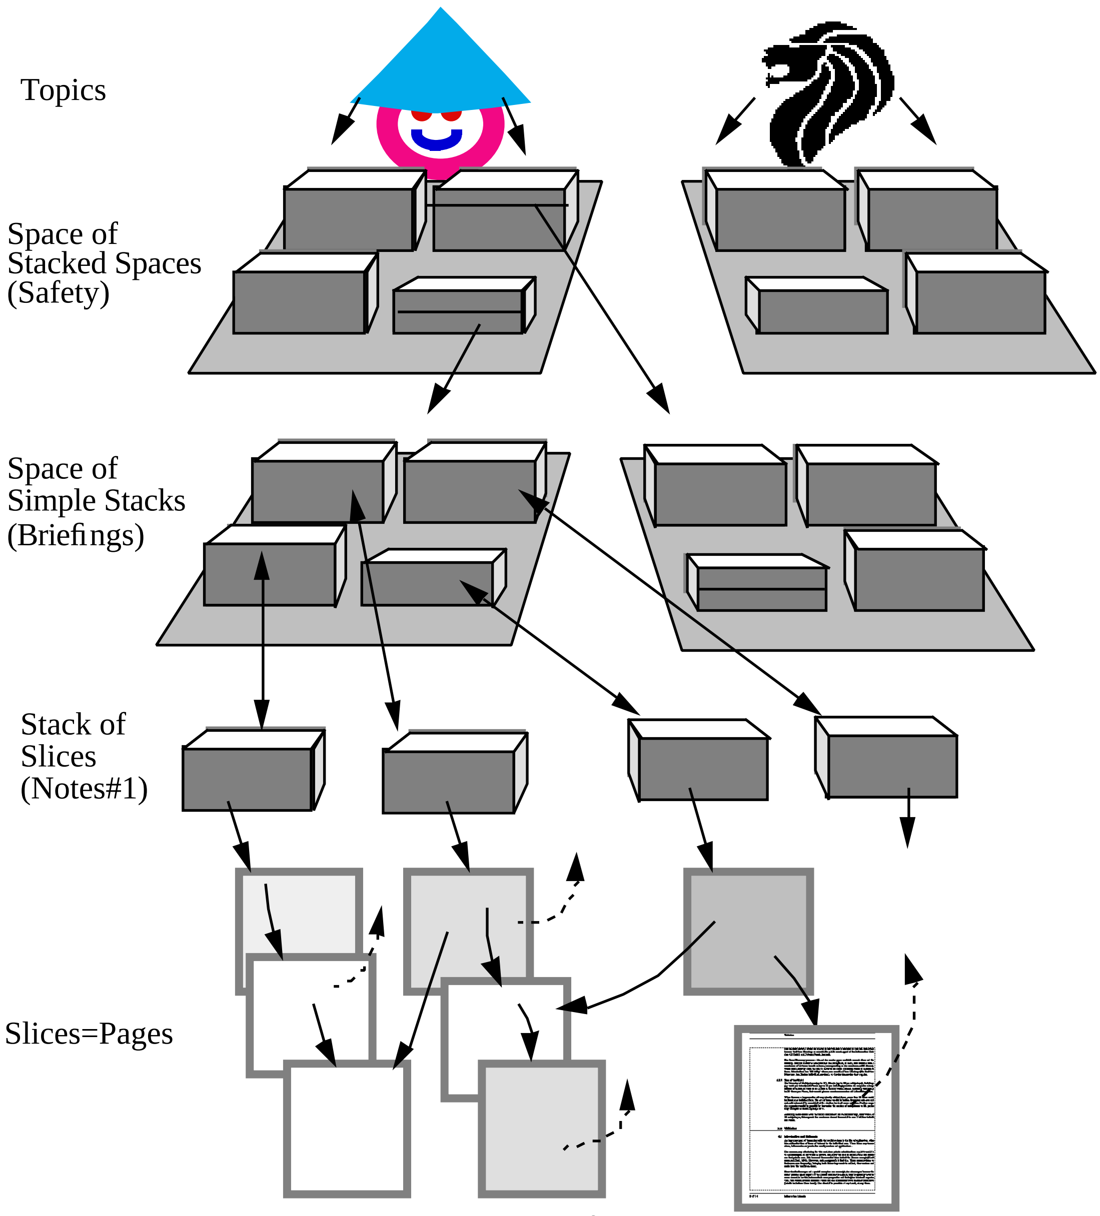

Figure 9 - Entities in the InfraSpace World

The major difference between InfraSpace and the general (but personal) exploration environment, StackSpace, is that the information is supplied and maintained by administrators in InfraSpace, whereas the user constructs and maintains his StackSpace world. The impact of this is that users have less flexibility in dealing with InfraSpace, unless they transfer materials to their own information space, because InfraSpace is a public information place. In contrast, StackSpace is an example of how a personal information space might operate.

Users operate within the context of the space they have entered. On entering a space, the set of stacks becomes visible with the topic marker forming a backdrop to the scene. The same background is maintained when dealing with individual stacks and slices. To change context, users need only raise their viewpoint until the other entities at a particular level become visible. With a high altitude viewpoint, all the topic markers will be visible. This functionality would be duplicated with a choice of topic menu. InfraSpace makes heavy use of transparency to allow users to ‘page’ through stacks to focus on the item of interest. Important and new items should be marked as such by the administrators. New items can be marked with flashing or otherwise animated markers. Size of items on the ground (“footprint”) is used to indicate their importance, perhaps simply reflecting how many staff have consulted that information. In contrast, height of a stack indicates the number of items of information it contains. So a tall, thin stack contains much information but has seldom been accessed, whereas a wide, low stack has been accessed often but contains little information.

Public passages (for example, walkways) between entities should be allowed to develop with use, as footpaths on a newly built estate. In this sense, the InfraSpace world begins life like an unused, planned city that then develops according to use (Lynch, 1960). The end product will be an information service whose structure, like that of a city, is a combination of planning and use - what is actually useful in practice (see Waterworth, 1996).

5. Discussion and Conclusions

Spatial interface worlds are of potential relevance to a wide range of information services and products. The basic idea of representing information as spatially navigable objects can be applied to the design of user access mechanisms to several types of large information/communication spaces (cyberspace), such as networked information services on the Web, video-on-demand, and other on-line services such as shopping, banking, electronic mail, etc.

The main focus of the paper was user exploration of the World Wide Web and how the design of ‘personal spaces’ could get around some outstanding problems such as providing true multi-threaded navigation support, the back-forward problem, and the need of users to casually organise, reorganise, filter and communicate information. A design for a users’ spatial Web environment which addresses many of these issues - StackSpace - was presented. The model was described in some detail at the conceptual level, relating the overall metaphor of personal spaces for exploration to specific topics such as multi-threaded navigation, chronology, currency, customisation and communication. StackSpace is essentially an environment for information explorers rather than providers. To investigate the generality of the organisational model behind StackSpace, it was applied to the problem of information provision. InfraSpace used the same elements as StackSpace, but attempted to apply them to the design of a spatial version of an existing information service - a public place.

There is, of course, potential for further development. Activities such as viewing and working with others, in a variety of ways and with varying constraints, would benefit from detailed study. The work could also be extended to cover more services, including services of different types. The exercise of designing the two types of application - a public place for information provision, and personal spaces for exploration - with the same type of spatial model is an interesting one, but the comparison between the two activities requires further study. In particular, a more detailed comparison of the differing requirements of information providers, as compared with information explorers, is an important topic for the future. So far, the only evaluations of these concepts have been ‘thought experiments’. These reservations notwithstanding, tentative conclusions are drawn below.

The tools and environments proposed in this paper are more suitable for information exploration, particularly in relation to navigating the Web, than for the design of specific services that provide information. However, an emphasis on accessing information is a useful basis for considering designs that purport to support the provision of information for user access. The model appears less successful for information provision than exploration, because it is probably less generally applicable to the former. Although the relatively simple structure seems to apply quite well to InfraNet, it is not clear that it would seem as reasonable if the information provided were more complex.

Personal Spaces, as embodied in StackSpace, are contrasted with public places such as InfraSpace. The design of public places, which provide information, is (paradoxically) less amenable to generic design solutions than is the design of personal spaces for exploring information. In the former case, things must be organised in ways that are understandable and useful to most people wishing to explore the information located there, but much depends on the nature and complexity of the information provided. A particular benefit of personal spaces, on the other hand, is that they use general tools and strategies to address the need for customisation by users of their interface to networked information systems. These tools and strategies are direct and should be readily understood by a broad class of users. They also provide mechanisms with potential utility in communicating and working collaboratively with information.

It may be that to improve the status quo of information provision substantially, providers will need to address the needs of explorers to a much greater extent than is currently the case. One approach would be for providers of information also to provide an environment and tools for exploration of the information provided. However, rather than encountering different exploration tools at different public places, each public place should be reasonably intelligible in its own right. Conversely, indviduals will wish to retain familiar facilities of their own personal spaces in their encounters with different information providers.

The on-line community as a whole is unlikely to tackle the difficult (and perhaps impossible) task of designing a single, unified cyberspace, despite the ever increasing importance of cyberspace for information access and communication. As Bruce Tognazzini points out (Fuller and Pemberton, 1995), tackling this task would require amiable and continuing cooperation, and is something which people seem determined to avoid (but see Waterworth, 1996). Given this state of affairs, the provision of tools and environments for information exploration, based on the notions of personal spaces and manipulable objects, is likely to be a key direction for the future development of cyberspace.

What users need, in trying to make sense of a vast collection of information like the Web, are the tools necessary to impose structure on what is, otherwise, an overwhelmingly amorphous mass. People look for and create structure and, it could be argued, all structure is spatial from the human point of view. The meeting of an abstract structure and a human being making sense of that structure does not result in the concept of using three-dimensional space as a metaphor. It results in real, three-dimensional, psychological space that can be simulated for the purposes of people using computers to explore information. Personal Spaces make sense.

6. Acknowledgements

This work was carried out as part of BT’s Short Term Research Fellowship Scheme.

Useful comments on earlier versions of the paper were provided by Thomas Ahlmark, Bob Benton, Martin Crossley, Victor Kaptelinin, Kristo Ivanov and especially Torsten Nilsson.

7. References

Andrews, K and Pichler, M (1994) - ‘Hooking up 3-D space: three-dimensional models as fully-fledged hypermedia documents’. In Proceedings of MHVR’94 East-West Conference on Multimedia, Hypermedia and Virtual Reality, Moscow, September 1994.

Brown, M H and Shillner, R A (1995) - ‘A New Paradigm for Browsing the Web’. In Conference Companion, CHI’95 Conference on Human Factors in Computing Systems (Denver, May 1995). New York: ACM.

Clancey, W J (1989) - ‘The knowledge level reinterpreted: Modeling socio-technical systems’. International Journal of Intelligent Systems, 8 (1), 33-49.

Colby, G and Scholl, L (1991) - ‘Transparency and Blur as Selective Cues for Complex Visual Information’. Proceedings of ISOE, Volume 1460, 114-124.

Dieberger, A and Tromp, J G (1993) - ‘The Information City Project - a virtual reality user interface for navigation in information spaces'. In Proceedings of the Symposium Virtual Reality Vienna, December 1-3, 1993.

Fairchild, K. M., Poltrock, S. E. and Furnas, G. W. (1988) - 'SemNet: Three-dimensional graphic representation of large knowledge bases'. In R. Guidon (Ed.), 'Cognitive Science and its Applications for Human-Computer Interaction'. New Jersey: Erlbaum.

Fuller, R and Pemberton, S (1995) - ‘Deconstructing Tog: An Electronic Interview with Bruce Tognazzini’. SIGCHI Bulletin, 27, 2, 24-27, April 1995.

Genau, A and Kramer, A (1995) - ‘Translucent History’. In Conference Companion, CHI’95 Conference on Human Factors in Computing Systems (Denver, May 1995). New York: ACM.

Halasz, F G, Moran, T and Trigg, R (1987) - ‘Notecards in a Nutshell’. Proceedings of ACM CHI and GI’87 Conference on Human Factors in Computing Systems. New York: ACM Press.

Hill, W C, Hollan, J S, Wroblewski, D and McCandless,T (1992) - ‘Edit Wear and Read Wear’. Proceedings of ACM CHI’92 Conference on Human Factors in Computing Systems (Monterey, May 1992). New York: ACM Press.

Jones, W P (1989) - ‘As We May Think? - Psychological considerations in the design of a personal filing system’. In R. Guidon (Ed.), 'Cognitive Science and its Applications for Human-Computer Interaction'. New Jersey: Erlbaum.

Lieberman, H (1994) ‘Powers of Ten Thousand: Navigating in Large Information Spaces’. Proceedings of UIST’94 Symposium on User Interface Software and Technology (Marina del Rey, November 1994). New York: ACM.

Lynch, K (1960) - ‘The Image of the City’. Cambridge, Mass: MIT Press.

Mander, R, Salomon, G and Wong Y (1992) - ‘A‘Pile’ Metaphor for Supporting Casual Organisation of Information’. Proceedings of ACM CHI’92 Conference on Human Factors in Computing Systems (Monterey, May 1992). New York: ACM Press.

Mukherjea, S, Foley, J D and Hudson, S (1995) - ‘Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views’. Proceedings of CHI’95 Conference on Human Factors in Computing Systems (Denver, May 1995). New York: ACM.

Musil, S and Pigel, G (1993) - ‘Virgets: Elements for Building 3D User Interfaces’. In Proceedings of the Symposium Virtual Reality Vienna, December 1993. Available as TR 93/13, Vienna User Interface Group, Lenaugasse 2/8, A-1080 Vienna.

Norman, D. A. (1993) - 'Things That Make Us Smart'. Reading, Mass.: Addison-Wesley.

Olsen, D R and Bialystok, E (1983) - ‘Spatial Cognition’. Hillsdale, NJ: Erlbaum.

Robertson, G G, Card, S K and Mackinlay, J D (1993) - ‘Information Visualization using 3D Interactive Animation’. Communications of the ACM, 36, 57- 71.

Robertson, G. G., Mackinlay, J. D. and Card, S. K. (1991) - ‘Cone Trees: Animated 3D visualizations of hierarchical information’. In Proceedings of CHI'91: Conference on Human Factors in Computing Systems (New Orleans, April 1991). New York: ACM.

Schön, D A (1989) - ‘Educating the Reflective Practitioner’, The Jossey-Bass Higher Education Series. San Francisco: Jossey-Bass.

Silvers, R (1995) - ‘Livemap: A System for Viewing Multiple Transparent and Time-Varying Planes in Three Dimensional Space’. In Conference Companion, CHI’95 Conference on Human Factors in Computing Systems (Denver, May 1995). New York: ACM.

Strong, G W (1993) - ‘Task Structure and Choice of Navigational Aid’. Proceedings of HCI International 93, Conference on Human Factors in Information Systems.

Waterworth, J A (1995a) - ‘Viewing Others and Others’ Views: Presence and Concealment in Shared Hyperspace’. Workshop on Social Contexts of Hypermedia, 16-17 February 1995, Umeå University, Sweden.

Waterworth, J A (1995b) - ‘HCI Design as Sensory Ergonomics: Creating Synaesthetic Media’. Proceedings of IRIS 18 Conference, Denmark, August 1995. B Dahlbom, F Kämmerer, F Ljungberg, J Stage and C Sorensen (eds), Gothenburg Studies in Informatics, Report 7, 1995.

Waterworth, J A (1996) - 'A Pattern of Islands: Exploring Public Information Space in a Private Vehicle'. In Brusilovsky, P, Kommers, P and Streitz, N (eds) 'Multimedia, Hypermedia and Virtual Reality'. Berlin: Springer Verlag Lecture Notes in Computer Science.

Waterworth, J A and Chignell, M H (1991) - ‘A Model of Information Exploration’. Hypermedia, 3 (1), 35-58.

Waterworth, J A and Singh, G (1994) - ‘Information Islands: Private Views of Public Places’. In Proceedings of MHVR'94, East-West Conference on Multimedia, Hypermedia and Virtual Reality. Moscow, September 14-16, 1994.

Yates, F. A. (1984) - 'The Art of Memory'. London: Routledge and Kegan Paul. First published in 1966.

