

The Knowledge Garden

Martin Crossley, John Davies,
Andrew McGrath and Marek Rejman-Greene
Knowledge Management Research
BT Laboratories
Martlesham Heath,
Ipswich, IP5 3RE, UK,
martin.crossley@bt.com
john.nj.davies@bt.com
andrew.mcgrath@bt.com
marek.rejman-greene@bt.com

"Excellently observed," answered Candide; "but let us cultivate our garden."
--Voltaire, *Candide*

Abstract

This paper describes the Knowledge Garden, a collaborative 3-D information visualisation tool. Since humans are instinctively attuned to a 3-D world, 3-D interactive multimedia information sources offer the possibility of an intuitive and natural way to explore complex ideas, concepts and information. The Knowledge Garden aims to provide an environment where users can meet colleagues and share relevant information among a community of interest. Information is seen as an organic resource that changes over time, and is represented as plants in a shared three-dimensional Knowledge Garden. Plants are grown from the results of an automatic clustering process whereby Internet resources containing related information are grouped together. We go onto describe proposed extensions to the Garden which would support and encourage the creation of communities of interest and the sharing of tacit knowledge.

1. Introduction

With the exponential growth of information widely accessible in modern society, knowledge management increasingly relies upon the ability to select information relevant to our needs efficiently, and the ability to manage emerging structures of knowledge.

As noted in [Gol98], geographically dispersed organisations particularly need knowledge management and organisational memory tools that encourage users to understand each other's ever-changing contextual knowledge and foster collaboration while capturing,

representing and interpreting the knowledge resources of their organisations.

Information filtering traditionally focuses on the relationship between a given query (or user profile) and information at hand. On the other hand, an exploitation of interrelationships among selected information may put otherwise isolated information into a meaningful context. The implicit structures revealed should help users to use and manage information more efficiently [Shi95].

The Answer Garden [Ack90] describes a system which provides a database of answers which grows organically as new questions arise and are answered.

This paper describes the Knowledge Garden, a collaborative 3-D information visualisation tool for finding, retrieving and sharing information. Since humans are instinctively attuned to a 3-D world, 3-D interactive multimedia information sources offer the possibility of an intuitive and natural way to explore complex ideas, concepts and information. The Knowledge Garden aims to provide an environment wherein users can meet and share relevant information.

The Knowledge Garden is one of a number of Knowledge Management tools that have been developed by the Knowledge Management Research team at BT Labs [BTL98b]. The tools enable users to find, summarise and share information within their working group, thus building up a body of knowledge through mutual contributions. Finding relevant information is not sufficient; finding who knows it, and sharing it with the relevant people is also crucial. As Samuel Johnson so elegantly put it over 200 years ago:

"Knowledge is of two kinds. We know a subject ourselves, or we know where we can find information upon it." -- Samuel Johnson, 1775

According to a recent Reuters report [Reu97], over 60% of managers believe that information overload is present in their own workplace. Moreover, 80% of managers expect the situation to get worse. An information famine has become an infoglut.

The copyright of this paper belongs to British Telecommunications plc. Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage.

Proc. of the 2nd Int. Conf. on Practical Aspects of Knowledge Management (PAKM98)
Basel, Switzerland, 29-30 Oct. 1998, (U. Reimer, ed.)

<http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-13/>

Information has historically been naturally self-managing in that it has been difficult to manipulate and difficult to disseminate. With the advent of the WWW, however, people have access to too much information, and too much has become worse than too little. Radically different strategies are required for managing both the information and the interaction of users with the information. As the information is placed in context and made available to those to whom it is relevant, so it is transformed into knowledge.

In the physical world our locality helps to control the information we access and the number of people we interact with. In the online world the relationships between people and information are difficult to manage. The 3-dimensional Knowledge Garden is an attempt to provide users with access to both relevant information and people. The garden is built using the WWW-based Virtual Reality Modeling Language (VRML), and hence is easily accessible from a WWW browser.

In the Knowledge Garden, information is imagined as *organic*, slowly changing, sometimes lasting a long time. It is a place for information that is important to the user, allowing it to flourish or wither and die.

"A garden...is a place that offers a sense of separation from the outside world. That separation must be created by making a threshold – whether real or implied. This threshold creates the ability to leave one world and enter another. That is the fundamental function of garden: to allow a person the psychological space to dream, think, rest, or disengage from the world... It is meant to describe another type of psychological space, to help you touch real self." [Sch97]

Section 2 of the paper describes 3D user interfaces, their uses and characteristics, and how they can be applied to information retrieval.

Section 3 introduces the Knowledge Garden, and describes how information and people are represented.

Section 4 describes how the concepts within Knowledge Garden can be extended to provide other tools for Knowledge Management, including its use for sharing tacit knowledge within communities.

2. 3D User Interfaces

2.1 3D and the Information Age

Three basic hypotheses regarding information provide the motivation for the work described here:

- In the Information Age the production of information is easy; it is the ability to find and gain access to relevant information that is the key to success;

- Since all data held in electronic form is digital, an iconic representation can be used to characterise attributes and relationships;
- Browsing and searching the WWW and the Internet can be a frustrating experience, particularly for inexperienced users [Dav97]. Users can feel lost, disorientated and overwhelmed by long lists resulting from a simple search. Information representation within the user interface using both 2D and 3D can reduce the time it takes to access information, make sense of it and enhance creative thinking.

2.2 Characteristics of 3D User Interfaces

Three dimensional user interfaces are now a common experience, particularly in the entertainment arena. Over the last few years there have been a number of significant technical advances; these include the proliferation of powerful desktop computers, the ability to generate complex 3D graphics and the increased connectivity to information and communication networks (particularly the WWW). This has opened the way to develop more meaningful ways of representing, and therefore communicating, information.

Waterworth [Wat96] highlights some of the advantages of 3D User Interfaces (UIs). Spatial UIs are beneficial due to the natural human capacities for dealing with objects in space, behaving in space and attaching meaning to spatial dimensions, and the intrinsic memorability of spatially arranged information. Spatial representation of information capitalises on innate human capabilities for dealing with space and for precise sensori-motor control. We look at some of these advantages in more detail below.

Automatic behaviour - Sensori-motor skills, with practice, can become automatic, like driving a car. They become largely unconscious, freeing attention so the user can focus on the reason for interacting, not on the process of interacting itself. However, automatic behaviour makes it harder to recall consciously. This could become a problem when navigating 3D spatial interfaces. Therefore, provision needs to be made for memory aids, signposts, landmarks and distinctive features within the interface which can help remind the user and locate him/her in space. Memory is essentially asynchronous communication with oneself. It is therefore important that the design of the spatial environment takes account of these human characteristics by providing references to orientate the user.

Perceptual understanding - Acquiring sensori-motor skills involves knowing how to go about doing something, rather than only knowing what is needed for the task [Cla89]. These skills, and interfaces designed to exploit them, tend to foster an exploratory approach to content once the basic style of interaction has been mastered. This therefore encourages 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. The key issue here for the interface designer is attaching meaning to the spatial arrangement of information within a 3D user interface.

Semantics of 3D UI's - Using three dimensions of space can enhance display, navigation and access of information. However, 3D instead of 2D doesn't guarantee good usability. Norman [Nor89] suggests spatial organisation of information only works when:

- There is a natural spatial mapping between (some, probably non-spatial, aspect of) the information objects and the spatial location;
- Desired items can be located in a minimum of attempts;
- The number of different items at any one location is small enough that they can be readily found;
- The amount of work and time required to try a location, scan its content, and then try another location is small.

It is the first of these points that is problematic but potentially powerful. It is essential to attach meaning in a sensible way so users can purposely explore these objects and interpret shapes and locations easily and naturally in terms of the concepts in which they are interested. This is an example of the use of metaphors.

Mnemonics (memory aids) - It is plausible that spatial interface models for information systems can themselves serve as mnemotechnic systems (artificial memory) for their users. Mnemonic systems generally work by associating to-be-remembered items with various locations within an imagined 3D world, like remembering the placing of a book in a room in a library. This is a classic memory technique, capitalising on imagery and motor-memory in location recollection.

Erickson [Eri95] also supports the potential of 3D Spatial User Interfaces (and indeed hints at their use in shared space environments which we consider in more detail later in this paper):

"3D seems to provide the possibility for representing many dimensions of information and meta-information in a compact and natural way. We believe that in the long run, information spaces will also be used as social spaces, and that 3D can serve as a natural framework for supporting social interaction. "

2.3 Information Retrieval in 3D

In the Information Retrieval (IR) field, interest has grown recently in trying to improve the effectiveness of information retrieval systems by improving the interaction between user and system and in particular by designing various types of graphical front ends (see, for example, [Cha93, Dav96]). In this view of IR, the role of the system is not to select documents similar to a user-supplied query but to organise and display

information about many documents in such way as to assist the user to select useful documents on his or her own. It is hoped the following advantages will accrue from such an approach:

- Users who are unfamiliar with the scope of a document collection will be able to familiarise themselves with it more easily;
- Users who are unable to express their information needs in 'traditional' query formats will be able to locate documents more effectively;
- Similarities between documents will be clearly presented which is not the case in a ranked list of documents matching a query.

Navigability in a purely symbolic domain has limits. The amount of 'depth' present in a subject before it exceeds human capacity for comprehension (and hence, navigation) is finite and relatively limited. Humans, however, have a sophisticated visual system and navigate in three dimensions; we are born to it and have a comprehensive ability to spatio-locate and spatio-organise. It thus seems reasonable to propose that the WWW should be extended, extending its navigation model from two dimensions into three.

2.4 Spatial Design for Information Retrieval in 3D

One of the most powerful features of 3D graphics is that if the graphical quality is sufficiently high then the interaction is sensually analogue. Instead of folders which are either open or closed a space allows information to be displayed in an infinitely variable granularity. With the increase in the power of agents able to synthesise and summarise information in a number of ways comes an increase in spatial metaphors allowing this information to be displayed in the same information space. This means the user can move amongst information, changing the depth of the view without leaving the original context.

One of the main drawbacks when exploring information in 3D is that it is not always straightforward to manoeuvre in a 3D space using a mouse. This is often exacerbated by poor spatial design. This drawback can however be overcome by thoughtful design of the world. For instance in the Knowledge Garden, the information finding agent which generates the plants (3D representations of information) also generates appropriate viewpoints from which to view those plants. Viewpoints, when selected, automatically fly the user to a predetermined position and orientation in the space. One can imagine combining these viewpoints with a customer's profile to generate a guided tour of the most appropriate information within a space. This also has the advantage opening up the possibility of serendipitous information finding: for example, while being taken to see something an agent has identified as being important to the user based on their profile, the user may fly past some information which catches their eye. This is a

specific example of a more general advantage of 3D representations which users navigate: with careful design of the space, the journey around the space can itself have meaning. We shall return to this point in section 4, where we consider extending the garden by exploiting information about social navigation (where do users go? Who do they meet? etc).

3. Information Representation in the Garden

The objective of the shared Knowledge Garden is to provide unified 2D/3D shared, dynamic and intelligent information interfaces and representations. The main focus of the work is user exploration and sharing of WWW-based information resources through the user interface. The Knowledge Garden should thus:

- Generate an easy-to-use, engaging, unified information interface;
- Build live, self-updating information places;
- Illustrate collaborative, shared access to information places;
- Provide communication (audio, video whiteboards etc.) for shared place User Interfaces.

The garden metaphor was chosen for a number of reasons:

- The concept of a garden is highly cross-cultural (although of course the particular implementations of gardens in different cultures vary widely);
- The metaphor reflects the organic, changing nature of information which it was hoped to reflect in the space;
- Gardens are generally relaxing, non-threatening environments, where the user can feel comfortable.

The project has built on previous work and extended the spatial metaphor of previous information landscapes and shared spaces as a framework for the user interface. Previous work includes BT's Portal and The Mirror [Wal96, Wal97, BTL98a].

3.1 Jasper Shared Information Store

A major consequence of the introduction of the WWW to the Internet has been an increase in the usability of the information within it. This is directly attributable to the navigability of the information: in other words, the WWW is useful (and will be used) to the degree it is capable of conforming to requests made of it. The WWW has added a universal structure to the information within it; through the WWW, over 300 million WWW pages [Sci98] can be treated as a single, unified data source, and all of the information can be treated as a single, albeit highly complex, document.

Given the immense amount of information available on the WWW, when information of interest is found, it is preferable to avoid copying entire documents from the original location to a local server. The usual bookmark facilities provided by browsers are however limited, typically allowing the local storage only of the title of a WWW page and its Universal Resource Locator (URL). The Jasper information agent [Dav95] is an attempt to address this problem. Users can 'store' pages of interest in Jasper. Jasper extracts the important details from these documents, including the summary and keywords, as well as the document title, its URL, and the date and time of access. This information is maintained locally and has several purposes. The summary gives a user an idea of the content of the document without the need to retrieve the remote pages. A quick link to the full remote information is also provided. Jasper locally indexes the pages that have been stored. Over time, a rich set of pages is built up by Jasper from different users' entries.

When a user stores a page in the Jasper system, that user's agent automatically informs other users with matching profiles, by e-mail. A reference is also added to every users 'What's New?' page along with a predicted interest rating. Furthermore, if the document does not match the profile of the user who stored it, then Jasper prompts that user to update their profile accordingly and suggests suitable additional keywords. In this way Jasper learns more about the user's interests as the system is used.

3.2 Information Representation within the Garden

Within shared environments the ability to generate WWW based information representations is key. Agents like Jasper, described above, are used to provide shared information resources and the ability to generate group or personal based searches and representations, providing another layer of information support inside the shared 3D spatial interface.

The pages stored in the Jasper system are exploited by making part of the 'Knowledge Garden' a 3D representation of the Jasper store, through and over which users can fly and wherein information is arranged in a logical, intuitive and accessible way. This arrangement of information is achieved by using the meta-information in the Jasper store and results from an automatic clustering process which groups together URLs (WWW pages) containing related information. The resulting cluster structures of the Jasper store are thus used to create a novel 3D front end onto the Jasper system, using VRML. This clustering work is described in more detail in [Dav96].

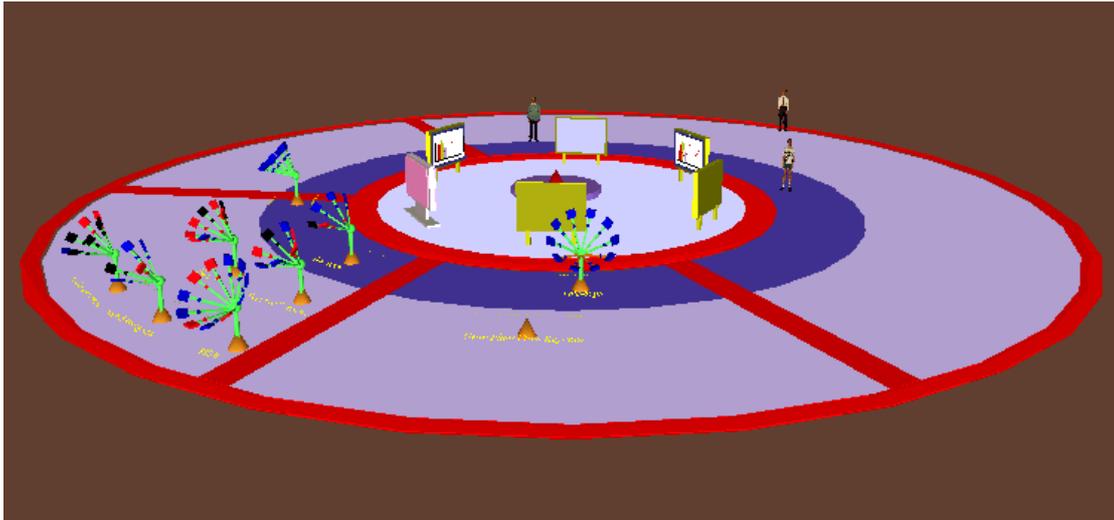


Figure 1: Shared Knowledge Garden

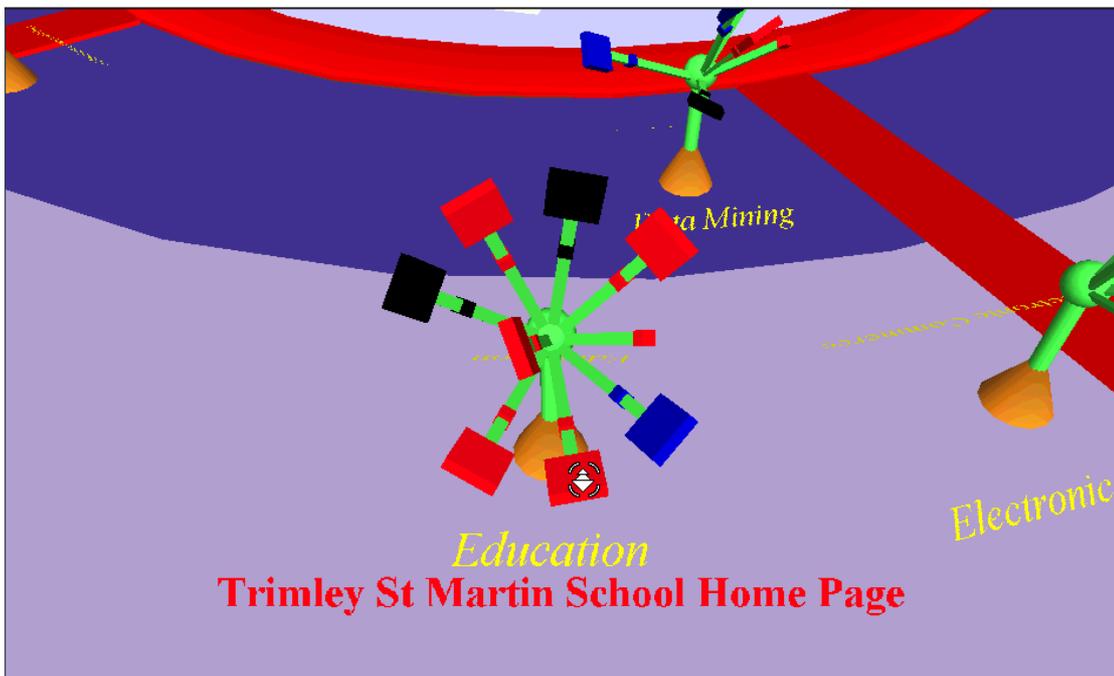


Figure 2: Knowledge Garden Document Cluster

Figure 1 is a screenshot from a working prototype shared knowledge garden. The central circle shown in is a shared area where avatars can meet and communicate via a number of media including text, audio and video. Shared electronic whiteboards are also provided. The outer sectors are different areas of the Knowledge Garden. Each sector represents a different information source. As the user navigates towards a particular sector, more detail becomes visible, as shown in Figure 2, where we can see clustered groups of documents (Data Mining, Distributed Computing, and so on). Each stalk represents a WWW

document, with a coloured 'flower' at its end. The colour coding of the flowers on the stalks represents the status of the particular document:

- Red indicates that a document has been updated since it was stored in Jasper;
- Blue indicates that the document has not changed since it was stored in Jasper;
- Black indicates that the link is 'dead' (that is, the URL associated with this document is no longer valid).

The smaller ‘buds’ part way up the document stalk represent a locally held summary of the information, as generated by Jasper. A heads-up display is used to show the title of the document as the mouse passes over the flower or bud.

When a document is selected, the associated URL is loaded into a separate 2D browser window. The selected stem then starts to wave, analogous to touching stems on a real plant as one walks through a real garden. The wave period decreases over time.

The user is able to prune unwanted plant stems; this enables the user to remove dead flowers, or areas that are no longer of interest. The pruning can either fully remove the stem, or just remove the flower, leaving the summary bud. Additionally, new plants can be grown from searches. Of particular note is that any changes made to the garden are persistent (that is, changes made in the virtual world persist even after the users leaves the world, and remain for future users).

At some point, it makes sense to drop down from a 3D to a 2D view of the information and by clicking on the flowers or buds, the user is shown the WWW page of the original information or the summary respectively.

The human eye is adept at picking up small movements, and can quickly spot dominant colours. Because of this, it is possible for the user to quickly view the garden, and pick out interesting aspects. For instance, an area of the garden which seems to contain a lot of red together with a lot of movement indicates that the information is new and has been accessed recently by other team members. Thus ‘hot news’ is shared among the community.

3.3 People availability within the Garden

Within the shared environment, users are able to establish the ‘availability’ of their fellow team members via their representational avatars within the people section of the garden.

Avatars are drawn for the team members, positioned dependent upon their perceived availability. Those who are currently more available are moved towards the centre of the garden, those less available towards the periphery. Their availability is calculated in real time, based on a number of factors, such as:

- Last time their mouse was moved;
- Last time a key on their keyboard was pressed;
- Number of unread email messages.

Additionally, team members can have a number of different avatar representations, dependent for instance on whether they are currently in the office or working from home.

3.4 Information Sharing within the Garden

By making the Knowledge Garden multi-user, users can not only access the rich variety of information resources within the Garden, but they can collaborate with other users across the Internet. When avatars meet in

cyberspace, they can communicate via a rich variety of media. The communications channels are set up automatically within the Garden. These include:

- Text;
- Speech (both IP and PSTN telephony);
- Audio/video (via CUSeeMe);
- White board.

Taking the gardening metaphor further, users can take ‘cuttings’ from information plants and grow their own copy within their personal environment. When a cutting is taken, the document represented by the stalk in question is analysed and a query in the form of a set of key phrases is extracted from it. So, for example, if the document was about network-centric computing the query extracted might be:

{network-centric computing, Java, network computers, mobile code, thin clients}.

This query is then passed to a search engine which searches WWW for documents which match the query. The ten most relevant documents are then represented as stalks on a plant which grows in the 3D space.

Thus users endeavouring to research information on a particular topic not only have access to the wealth of links stored within Jasper and any other information sources represented in the garden, but can be automatically put in touch with (perhaps unknown) colleagues who have expertise in this field and possibly meet them in the shared environment.

3.5 Evaluation of the Knowledge Garden

An early version of the garden has been in use by project team members for some time. Although no formal evaluations have been carried out, features to which they have reacted particularly positively include:

- The ‘organic’ nature of the garden metaphor, which encourages the idea of information which will grow (and be pruned);
- The ability to visually represent a reasonably large collection of documents at different levels of detail;
- The meta-information about the information collection provided by clustering and colour coding of the documents, particularly when viewed at relatively long distance - examples of the information provided include the amount of information available on a given topic and clusters where many links have ‘died’ or need updating;
- The tendency to get ‘lost in hyperspace’ is reduced, since the user can stay in the same 3D space and yet have access to a range of information and meta-information at multiple levels of details;
- The use of automatically-generated pre-defined viewpoints within the Garden help navigation considerably.

4. Extending the Knowledge Garden

The value of the Knowledge Garden can be enhanced by integrating it with other sources of information. In its initial application, the Garden displays information in the form of pointers to resources on the Web. By focusing exclusively on this information and its first-order relationship to other similar items, such a system misses out on other valuable attributes of the knowledge transaction:

- The identity of the individual who alerted the system to the existence of the resource;
- Other resources which were accessed in the same Web session by the alerter;
- The creators of the resource;
- The identities of people who subsequently accessed the resource using the Garden, their comments on the resource, and the relative times at which these actions took place.

Further development of the tool will demonstrate the power of using these additional elements in identifying existing communities of interest within an organisation. A visual representation of the expertise across the organisation can thus be generated. By switching from the information view to the people view, the focus moves from what to who. The asynchronous interaction between suppliers and viewers of specific items of information in the Garden establishes a web of links between users of the space.

Individuals may of course be multiply represented if they have multiple interest areas, and the spread of an individual's contributions could be highlighted as the cursor moves above a specific instance of that individual's participation. By redrawing this web in different ways, it is possible to:

- Emphasise the clustering of interests in communities;
- Discover gatekeepers to expertise which is currently remote from the viewer's specialisms.

In this way, intermediaries are identified, through whom an enhanced relationship of trust can be established, thereby improving the likelihood of a more productive sharing of knowledge between parties who may never have met.

A large amount of the knowledge within an organisation may of course not be codified: it may be personal, context-specific and difficult to write down. Such knowledge is referred to as tacit knowledge [Pol66] and access to and sharing of this type of knowledge is a major concern of the knowledge management undertaking. When tacit knowledge is difficult to make explicit (codify), we need to find new ways of transmitting the knowledge through the organisation. Failure to do so can lead to loss of expertise when people leave, failure to benefit from the experience of others, needless duplication of a learning process, and so on.

The current shared space aspects of the Knowledge Garden are functional but limited. Further development

needs to be undertaken to improve the scalability, such as the inclusion of focus, nimbus and aura [Ben93]. The users profile could then potentially be used as a basis for the aura, rather than just the traditional distance-based aura. This would enable users to collaborate with others with similar interests, not just those in closest proximity.

One way in which the Knowledge Garden can encourage the sharing of tacit knowledge is by using its knowledge of the users within a community of interest to put people who would benefit from sharing their (tacit) knowledge in touch with one another automatically. In this way, the garden would encourage the exploitation of "weak ties" in an organisation. These relationships are those outside a worker's immediate group and have been found to be very important the process of sharing knowledge across organisation subunits in a multiunit organisation (see for example [Han98]).

In this way the Knowledge Garden, while not claiming to actually capture tacit knowledge, provides an environment which actively encourages the sharing of tacit knowledge, potentially by people who previously would not otherwise have been aware of each other's existence.

5. Concluding Remarks

We have described the Knowledge Garden, a WWW-based tool for sharing information via a natural and intuitive 3D interface. The benefits of the 3D approach were explored and the way in which the Garden uses various visual cues to highlight different attributes of information was discussed. The Garden is a shared space, and we described a set of utilities which have been provided to allow users to interact with one another while in the information space. Finally, ways in which the Knowledge Garden could be extended were discussed with particular emphasis on extending it from a simple information space to a space in which relationships between people could be teased out and exploited to encourage and facilitate knowledge sharing.

Acknowledgements

The authors gratefully acknowledge the contributions of the following to the development of the Knowledge Garden: Chin Weah Chia, Rob Taylor-Hendry, Amanda Oldroyd, Matt Shipley, Phil Clarke, Andrew Hockley, Lewis Collins, Paul Rea and Peter Maydell.

6. References

- [Ack90] M.S. Ackerman and T.W. Malone. Answer Garden: A Tool for Growing Organizational Memory. *Proceedings of ACM Conference on Office Information Systems*, 1990: 31-39.
- [Ben93] S.D. Benford and L.E. Fahlen. A Spatial Model of Interaction in Large Scale Virtual Environments. *Proc. Third European Conference On Computer Supported Co-operative Working*

(ECSCW '93), Milano, Italy, Kulwar Academic Publishers, pp. 109-124.

- [BTL98a] *BT Labs – Shared Spaces*, <http://virtualbusiness.labs.bt.com/SharedSpaces/> correct at 30/9/98
- [BTL98b] *Knowledge Management* http://www.innovate.bt.com/showcase/knowledge_management/ correct at 30/9/98
- [Cha93] M. Chalmers. Using a Landscape Metaphor to Represent a Corpus of Documents. in Frank, A U and Campari, I (Eds.) *Spatial Information Theory* Springer-Verlag, 1993.
- [Cla89] Clancy. The Knowledge level reinterpreted: Modelling socio-technical systems. *International Journal of Intelligent systems* 8 (1) 33-49 1989.
- [Dav95] N.J. Davies, R. Weeks and M.C. Revett. An Information Agent for WWW. *Proc. 4th Intl. Conf. On World Wide Web, Boston, USA*, December 1995. Also available at <http://www.w3.org/pub/Conferences/WWW4/Papers/180/>
- [Dav96] N.J. Davies, R. Weeks, M.C. Revett and A. McGrath. Using Clustering in a WWW Information Agent. *Proc. 18th BCS Information Retrieval Colloquium* British Computer Society, 1996. Also available at <http://www.labs.bt.com/jasper/html/jasclus2.htm>
- [Dav97] N.J. Davies and M.C. Revett. Networked Information Management. *BT Technology Journal* Vol 15 No 2 :194-208, April 1997.
- [Eri95] T. Erickson and M.P. McCahill. *A preliminary design for a 3D Spatial User Interface for Internet Gopher* 1995.
- [Gol98] R. Goldman-Segall and S. V. Rao. A Collaborative On-Line Digital Data Tool for Creating Living Narratives. in *Organizational Knowledge Systems*, presented at The 31st Hawaii International Conference on Systems Science, Hawaii, HA, 1998.
- [Han98] M. Hansen. *The role of weak ties in sharing knowledge across organization subunits*, Harvard Business School, Working Paper, 98-011, 1998.
- [Nor89] Norman. *Things that make us smart*. Reading, Mass, Addison-Wesley 1989.
- [Pol66] M. Polyani. *The Tacit Dimension*. Routledge and Paul, London, 1966.
- [Reu97] Reuters. *Glued to the screen*, Dec 1997. ISBN 0-901249-06-8.
- [Sch97] M. Schwartz. *Transfiguration of the commonplace*. Spacemakers Books, Washington DC. 1997 ISBN 1-888931-01-9.
- [Sci98] S. Lawrence and C.L. Giles. Searching the World Wide Web, *Science*, 280, p. 98, April 3, 1998. Also available at <http://www.neci.nj.nec.com/homepages/lawrence/websize.html>
- [Shi95] F.M. Shipman, C.C. Marshall and T.P. Moran. *Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information*, presented at CHI '95, 1995.
- [Wal96] G. Walker, J. Morphet, M. Fauth and P. Rea. Interactive Visualisation and Virtual environments on the Internet. *British Telecommunications Engineering*, Vol. 15, April 1996.
- [Wal97] G. Walker. The Mirror - reflections on Inhabited TV. *British Telecommunications Engineering*, Vol. 16, April 1997. Also available at http://virtualbusiness.labs.bt.com/msss/IBTE_Mirror/index.htm
- [Wat96] J. Waterworth. Personal spaces: 3D Spatial Worlds for Information Exploration. *3D and Multimedia on the Internet, WWW and Networks*. BCS International Conference. Bradford, UK 16-18 April 1996.