Journal of Universal Computer Science, vol. 6, no. 5(2000), 503-516
submitted: 21/7/99, accepted: 4/5/00, appeared: 28/5/00

Treeworld: A Conceptual Model for Large-Scale Hypermedia

D.B. Skillicorn
Department of Computing and Information Science
Queen's University, Kingston, Canada
skill@cs.queensu.ca

Abstract: Existing interfaces to large-scale hypermedia such as the world wide web have poor conceptual models and poor rendering of navigational and contextual information. New technologies that make it cheaper to use three-dimensional representations suggest the use of richer conceptual models. We discuss criteria for assessing more powerful conceptual models and design decisions that have to be made to exploit richer interfaces. The Treeworld model is suggested as one attractive example of such a model.

Keywords: conceptual model, navigation, world wide web, large-scale hypermedia, 3-d glasses, search, relevance structuring, hierarchy, focus + context, visualisation, teleportation, Treeworld.

Category: H.5 Information Interfaces and Presentation

1 Introduction

Most existing interfaces to large-scale hypermedia have limited conceptual models, with simple interfaces that implement them. In practice, this means that someone trying to navigate such structures has limited choices at any moment, has a limited view of the local context, and can extract only minimal meta-information.

The aridity of a conceptual model of hypermedia systems such as the world wide web arises from its limited link model and the wildly autonomous control of content and structure on which it is based. At the same time, interfaces are limited to flat screens of limited area that are required to render both content and structure in the same setting.

We suggest a much richer conceptual model, based on normal human navigation in a three-dimensional world, a world that can now be replicated at low cost for individual users via 3-d glasses. This opens up the possibility of separating hypermedia meta-operations from information processing operations, making navigation through structure an explicit operation. The approach is probably too computationally demanding to be feasible today, but this might be expected to change rapidly as processing power increases.

2 Using Large-Scale Hypermedia

There are five classes of activities that can be distinguished in large-scale hypermedia information systems such as the world wide web. They are:

- Find – search for something in a given context.
- Find again – recall something already encountered.
- Understand – absorb the content of a set of nodes, which is almost always significantly enhanced by awareness of the context in which they occur.
- Browse – wander through information structures in unplanned and serendipitous ways (this is sometimes called 'foraging').
- Interact – perceive and respond to the actions of others using the same information resources.

In hypermedia systems of moderate scale, the purpose for their existence and the implicit contract between content providers and content users is fairly explicit. For example, a help system exists for a well-defined purpose, primarily related to the first two activities above, and this guides many aspects of its design. In contrast, extremely large hypermedia systems, of which the world wide web is only one, have multiple reasons for existing and hence more diffuse relationships between content providers and content users. For example, the world wide web has many characteristics of a large billboard – content is provided in the hope that it will be noticed, but there is little control of who actually will encounter that content (and therefore few plausible assumptions about their characteristics). These diffuse relationships prevent the use of context from playing a major role, although it will clearly become more significant. Note that some context is already used – downloading a software patch from the web site of a major organisation is more likely than from an individual's home page, and there are many attempts to use brands or pseudobrands to give individual sites an imprimatur. The actions listed above have been expressed from the perspective of content users, but each has its implications for content providers as well (for example, in the presence of browsing, advertisers work to make it likely that their sites will be visited from sites that are popular based on their content – advertisers are parasitic in their approach).

Given these actions, it is natural to ask how easy a given hypermedia system makes them. The answer depends on a number of characteristics of such a system: content indexing, navigational techniques, model of arrangement of content, rendering of metacontent and so on. We concentrate on the aspects of navigation in its fullest sense, which has sometimes been called the *conceptual model* of the hypermedia system.

3 Conceptual Models

Conceptual models are the metaphors that a given information system presents to users. They have three components: the information itself, the way in which the information is organised, and the navigation paradigm by which new information can be discovered. The conceptual model determines what actions are possible and helpful at an given time.

Almost all of the conceptual models for large-scale information systems have an active role for the user; they convey, to some degree or other, the idea that a user moves to locations where data is to be found. This is a natural outgrowth of the physically distributed nature of large information systems. Some conceptual models (e.g. Deckscape [4, 5]) reflect more closely the physical reality: that all data is actually transferred to the user's own computer. Nevertheless, what might be called the 'travelling user' paradigm seems to have been useful. Note, however, that it is not user-centric, but rather location-centric, and so in fact organisation-centric. This may be a negative aspect of such metaphors; indeed arguably the poor mechanisms for bookmarking in most browsers are one consequence.

Three main kinds of conceptual models have been applied to large-scale information systems such as the world wide web.

The first is the default model of the web itself: nodes connected by hard-wired links inserted by authors, supplemented by search engines. Finding information, although a subject of frequent complaint, is reasonably effective. Finding information again is possible over short time scales, but changes in the raw data available to search engines and their own internal non-determinism seriously limit this. Browsing is constrained to compositions of paths envisaged by authors. The experience of finding and absorbing information is a solitary one, and context for any given node is limited to its URL which gives a small glimpse of one hierarchy in which it is embedded. The limitations of this conceptual model are graphically shown by work on modelling browsing. For example, Huberman *et al.* [8] show that typical user browsing behaviour can be accurately modelled as a sequence of moves that continue as long as sufficient prospect of value is provided by each new document. As a result, actual hits on web pages can be modelled using spreading activation based on document-to-document transition probabilities. In other words, a model based on random walks accurately models measured hit rates in the web. Clearly, users are unable to move around the web in ways that effectively satisfy their goals.

The second is the class of models that might be called relevance structuring, or perpetual search [16]. In such models, each node is connected on the fly to a list of successor nodes ordered by how much they are relevant to the content of the current node. The process may be initialised or primed by search terms. Finding is the basic operation of these models; and finding again is probably easier than in the web model, although again subject to the vagaries of the underlying engine. Browsing is natural in such models, but can only be done on the basis of content. Once again, the experience is a solitary one; and context is non-existent apart from the intellectual locality implied by relevance.

The third class of models are those that impose or infer meta-structure on nodes and present this higher level structure as a 'map' through which to navigate. There are a number of different alternatives depending on how this meta-structure is created. In the simplest case, it is simply the graph structure of explicit links, perhaps rendered to try and show certain structures to best advantage. An example of this are search engines that consider the documents returned by a search to be a subgraph, and which try to find sources and sinks in this subgraph. Such nodes tend to be more useful than their pure content would suggest – for example, a document with many links to other documents returned by the same search may be a tutorial. Potentially, this approach can be based on any kind of meta-structure which can be practically computed. (A major problem is that individual nodes lack identification of their internal meta-characteristics, so that it is hard to build on this weak foundation to infer relational meta-structure.)

Such models improve the ability to find information because they make locality visible. Thus finding can use the combination of getting close and then navigating to avoid the limitations of relevance and reliability. Finding again is similarly helped by the presence of higher-level structure into which information is fitted. Browsing is more flexible than in relevance structure models because adjacency is based on properties other than simple content. The rendering of an image of a meta-structure allows, in principle at least, the inclusion of the actions (and even the avatars) of others, so that some sense of collective action is possible. Provided that the meta-structures used are appropriate, context is omnipresent.

Models of the third class are of increasing interest as devices used for interacting with large information spaces become closer to virtual reality. Such devices make it easy and cheap to represent complex information visually, and hence to exploit meta-structure. For example, i-glasses (www.i-glasses.com) allow two- or three-dimensional presentation of a screen image in a pair of goggles for about $US500. The vision of the novelist William Gibson of a world-wide information system accessed visually is now cost-effective.

4 Issues for Meta-Structured Conceptual Models

A meta-structuring model must decide how to extract meta-information about nodes, how to model this information, and how to render the model so that the combination is as effective as possible.

The following properties should be considered in deciding whether or not a given meta-structured conceptual model is attractive:

- Is the meta-structure *visible* or *implicit*? Systems may either display meta-structures and allow them to be interacted with directly; or insert extra links to make the navigational choices of the meta-structure available (for example, Hyperwave [10]).
- Does the meta-structure make *finding* easy? For example, it should be possible to reach the appropriate part of the metastructure easily, and clear how to move from there to the precise node desired. At present most systems rely on ancillary mechanisms such as search engines for this functionality, but strongly hierarchical systems such as Hyperwave and Yahoo provide it directly.
- Does the meta-structure make *finding again* easy? For example, is the meta-structure static, and is the perspective on it repeatable over time? (See especially [19] for a discussion of the importance of this issue.) Very few systems address this functionality at all, mostly because the freedom to create and destroy content in present hypermedia systems is taken to be equivalent to license to do so arbitrarily.
- Does each node's *context* in the meta-structure correspond to at least the context assumed by its author when it was constructed? ([15]) This functionality is very clearly absent in the world wide web: a page found via a search engine is often incomprehensible because the author never envisaged it being accessed directly as a lone object.
- Does the meta-structure make *browsing* easy? For example, are reasonable sequelae of the current node all visible?
- Are the *actions of others* visible in the meta-structure? There are a wide range of possibilities here, from visible avatars of other users, to visibly denoting newly-created, newly-deleted or busy nodes. ([7])

- Does the meta-structure use 'screen' real estate well? Many web sites, for example, try to cram as much content into as few pixels as possible. Good use of screen real estate has to be much more aware of what is known about the human visual system and human attention to provide better *information* rather than dense content.
- Does the meta-structure use a *metaphor* that is appropriate and natural for humans? ([18])
- Where does the meta-structure come from? Is it based on information supplied (explicitly or implicitly) by the author, or is it generated from the node data or metadata (for example, using latent semantic indexing [3])? In either case, is the resulting structure *rational*? ([12, 16])
- Does the meta-structure *change* and, if so, over what time frame?
- Does the meta-structure make *effective use* of size, position, colour, and movement in rendering its information?

5 Existing Structured Conceptual Models

Several structured conceptual models exist:

- **Hierarchies.** Hierarchies arise naturally from the directory structure used in most modern file systems, and from the structure of most large organisations. Hierarchies are implicit structures even in the world wide web. Several systems have attempted to extract, render, and use this hierarchical structure. Systems such as Hyperwave (Hyper-G) [10] make hierarchies of collections a navigational alternative by adding up and down links to all pages served. Web portals such as Yahoo do the same thing 'manually' as an organising principle. Several systems also try to extract hierarchy information and render it in helpful ways. The hierarchical views approach uses this idea in a fundamental way [13, 14]; while many other systems use rendering techniques such as cone trees [17], tree maps [9], overlapped trees (Cheops) [2], and pyramids [1].
- **Focus plus context.** When a structure is too large to be presented completely with enough resolution to exploit local information, it is natural to use fisheye techniques to present a region in detail, with the remainder rendered in a way that preserves its general character. This has been called focus+context [11].
- **Physical world analogues.** Humans are used to navigating in the real, physical world. It is therefore natural to exploit this by making navigation in information systems resemble navigation in the real world. This idea is very general and works out in a number of ways. Some hypermedia systems allow nodes to be automatically organised into rings or tours in which each node has a defined predecessor and successor. Other systems treat clusters of data as buildings or rooms within buildings, with a navigational metaphor of movement from place to place. Libraries are a popular metaphor because we are used to searching and finding information in them. Some information systems have even adopted the networks of caves common in some adventure games as a natural way to structure information [6]. One of the biggest advantage of these approaches is that humans are accustomed to remembering spatial relationships once experienced, so that finding again may be easier in such systems.

- **Existing link structures with clustering.** Approaches that begin with some particular meta-structure in mind must generate or find it. When data already exists in some other shape, such as explicitly-connected web pages, it may be difficult to do this. Rendering meta-structure from arbitrary graph structures can still be useful, but there cannot be a predetermined form. Instead, systems render whatever structure they find, typically with some form of abstraction, perhaps so that clusters can be seen.

6 The Treeworld Conceptual Model

6.1 Basics

Buildings have attractive properties as natural maps of information spaces. Indeed, as far back as Greek oratory, the layouts of buildings were used to organise memory (although Greek buildings imposed a structure which is quite different from that of modern buildings, whose basic structure is hierarchical, and can often be closely mapped to organisational hierarchy). Buildings are not an ideal metaphor for hypermedia because they are over-constrained by the technology needed to construct them. For example, buildings have only a few entry points from outside, primarily at a single level, and they enclose their subspaces, so that moving to a low-level unit means traversing many layers of access space. Both these attributes do not fit well with the natural organisation of data.

However, consider what happens when a building is turned inside out. The result is a tree, not a computer scientist's tree but a naturalist's one, in which the main entrance is at the bottom of the trunk, each of the regions of the building are branches, and this structure replicates itself at smaller and smaller scales, with the leaves corresponding to single rooms.

Thus trees represent the natural hierarchical structure that is present in most information nodes (web sites) in the same way that buildings do. But they also display their leaves on the *outside* of the structure, making them natural targets for direct access. Representing an information node as a tree allows it to be accessed hierarchically and directly with equal ease and naturalness. In other words, trees have all of the structural and navigational affordances of buildings but they have the extra affordance of direct access to any node and they do not require the overheads of access structures (lobby, elevators, hallways).

In the Treeworld model, each web site is rendered as a tree using the directory structure of the files it contains. Thus a directory that contains five subdirectories is rendered as a node with five branches going upward from it to the nodes corresponding to these five directories. In the obvious way, a file named `index.html` is associated with the node corresponding to the directory in which it occurs. Other non-directory documents also form branches above the node corresponding to the directory in which they occur. Figure 1 illustrates the mapping between directory contents and the corresponding rendering. Browsing in a Treeworld setting, a user sees a collection of trees, and may move about either by direct, apparently-physical movement in the rendering, or by selecting the image of a node and moving directly to it. A user can therefore be in one of two states: 'inside' a document, viewing its contents, or 'outside' in the rendering of a location in the meta-structure representing a part of the hypermedia.

Each node in the structure has a natural navigational framework, with two canonical directions: rootward and leafward, with multiple choices in the leafward

```
home
 index.html
 summary.html
 products/
 prices/
 update.html
```

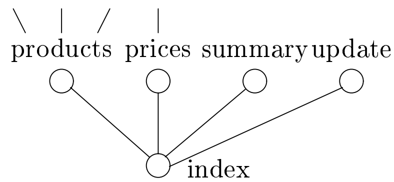

Figure 1: Correspondence between directory and rendering. The node labelled `index` corresponds to the directory `home`. The nodes `update` and `summary` are leaves.

direction. However, unlike the world wide web, the other navigational choice is not to follow a limited number of links determined by the current node's author, but to move (jump) directly to any visible node. This flexibility can be augmented by an authored set of choices that are represented, for example, by explicit (visible) links, or sets of nodes which flash to indicate their appropriateness as next to be visited.

Contextual information comes for free. Following the rootward and leafward links carries with it an implicit sense of 'movement' and hence accumulatively of position. Following arbitrary links causes contextual information to be displayed depending on the movement-rendering paradigm chosen, but a sense of location of a link target in the tree is a minimum context that seems unavoidable.

Trees are also a natural way in which humans are accustomed to perceive space. Exploiting 3-d space allows some trees to be near, and others to be occluded or distant while still presenting some visual clues of their presence and characteristics. The shapes of entire trees carries information about the branching structure of the information they contain, but cues such as colour and shape of both branches and leaves can be used to render further information.

Trees also have the advantage that they are fractal. It is thus essentially irrelevant, for example, if the user's starting point is a collection of trees or simply a branching point in a much larger tree (the World Wide Tree, even).

Trees have both an inside and an outside, allowing further freedom to label branches and leaves with their distance and orientation from the tree's centre. For example, common but specialised entry points can be placed on the outside of the tree, while their sequelae can be placed in a positions that are visible from the nodes on the trees periphery, but not necessarily from outside.

It is sometimes important to provide a sense of which information spaces are popular. This could be done by representing real-time accesses to individual pages, but this is computationally unappealing, and raises privacy concerns. It is plausible, however, to represent the same information in a way which avoids both of these problems. If a statistical model of accesses to each page is known, it can be used to render typical traffic patterns. Thus a user might see the movement of fictitious others to and from nodes of trees; and their overall motion faithfully captures the way nodes are actually being accessed.

To give some idea of how web sites appear when rendered in this style, we illustrate the rendering of two real web sites. Figure 2 shows the web site of a small company selling a technical consumer product, while Figure 3 shows the web site of a large non-profit technical organisation. Only the directory structure is shown, with numbers indicating the number of content pages above a node in some places. In the second figure, no unlabelled node has more than 10 content pages above it. The two web sites contain about the same number of total documents. However, the first is characterised by being of low height, having a large branching factor, and being relatively homogeneous. Its relatively simple structure is immediately evident in the rendering. In contrast, the second web site is taller, with a much smaller branching factor, and much less regular structure. This reflects the much greater complexity of the second site. It is not shown in the rendering, but the second site also has more frequent explicit links between different branches.

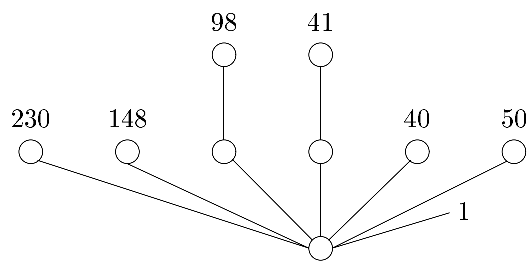

Figure 2: Web site of a small technical company

6.2 Navigation

As we have seen, there are two modes of moving from one ‘location’ to another. If the user is located at a node in a tree, then the rootward/leafward mode is the natural way to move – corresponding to climbing around the tree’s branches. However, even in this situation it is natural to want to move to different nodes within the tree directly, and to move to nodes in different trees. The implementation of this is obvious – a user moves to a new node by clicking on its representation in the rendering. However, this immediately raises the issue of what a user’s *field of view* is from moment to moment.

It is perhaps natural to begin with a view of a set of trees as if the user were standing in a wood. However, once the user is conceptually at a position inside a tree, it is less clear what view should be rendered. An ‘outside’ view would be of the same general sort as the initial view. However, it may be necessary to allow an internal view (of the ‘inside’ of the tree), or even to allow generic panning of the view.

A move to a new node includes two possibilities: opening the new node and placing the user ‘inside’ it, i.e. viewing its content; or rendering the view from the new node (i.e. remaining ‘outside’ it). The question then becomes how to represent the visual field during the transition. If the visual representation is an approximation of virtual reality, then effects such as vertigo must be taken into account. Using a motion to which users are accustomed in the real world might be expected to help.

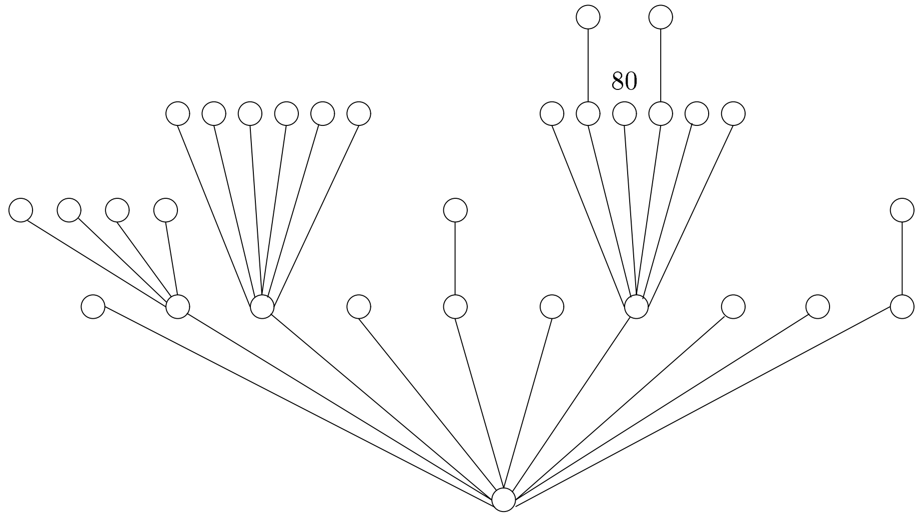

Figure 3: Web site of a non-profit technical organisation

There are several plausible possibilities:

- The motion is represented as a pendulum swing (i.e. like Tarzan).
- The motion is represented as sliding down or up a straight line between the current and new nodes.
- The motion is the parabolic curve of flight if fired from a cannon at the current node and landing at the new node.
- The motion is represented as as jump from the current node to the new node.
- The motion is represented as a helicopter ride between the two nodes.

A third, intrinsically different, form of navigation is required: *teleportation*, allowing instantaneous movement to a completely different part of the metastructure. This might be allowed freely, as hard-wired links are permitted anywhere in the world wide web today. However, there would seem to be advantages, both for simplicity and flexibility, in placing teleporters at fixed, known, locations such as the base of trees. Having a static network of teleporters makes it possible to provide directory services in a predictable, repeatable way. It has the further advantage that the number of entry points to a region of the virtual space is limited, so that resources can be spent precomputing the view at each one.

6.3 Larger Scale Structure

We have so far described a setting in which a user is placed in front of a collection of trees. There are a number of important issues to do with what this local neighbourhood or 'wood' should contain, and how such local neighbourhoods change.

The first major issue is the global view of the hypermedia system. There are two fundamentally different choices:

1. There is a single global structure to the data which all users see in the same way. This is essentially the structure of present world wide web, although it is somewhat obscured by the sheer rate of change of web pages, and by the autonomous ability to create new content.
2. The global structure is user-centric and different users perceive different arrangements of the data. This is the approach of the Hyperwave system, and other hypermedia systems that use link bases. The links visible, and hence the navigation possibilities, depend on each individual user's status.

This basic choice affects how a number of subsidiary choices are made.

A single global structure. A single global structure is the norm when the hypermedia system is under the control of a single organisation or unit. It is emphatically not the norm in the current world wide web. But even here there are elements of global structure, for example in the way that domain names are administered. The chaotic churn of web pages makes so many problems so hard that there is some pressure towards a more controlled structure, although what it might be is hard to predict. In any case, the only essential requirement for our use of a single global structure is that it is common rather than generated, at least in part, for each user.

The first issue is access: when users 'enter' the hypermedia system, where are they and what do they see? There are three kinds of solutions. In the first, the global structure has a set of common portals which are used by all users. An individual user can select which of these to use, but access cannot be customised beyond this. In the second, there are common portals but they are customised by user. The web shows a trend towards a limited form of this kind of access, in which internet service providers attempt to customise the initial page each user sees. In the third, each user can select a point in the global structure to act as his or her personal portal. The personal portal could be a particular point in the existing global structure, or it could be an array of teleporters linking to multiple points. A rudimentary form of this exists in the web in the concept of a home page which is loaded on browser initialisation.

The second issue is navigation outside the viewable destinations at any given location. There are two possibilities: common teleporters included as part of the global structure, or individual teleporters created by users. (Note that this corresponds to authored links versus user-created links, e.g. annotations, in existing hypermedia.)

User-centric structure. When the structure rendered depends on the individual user, issues of access and movement become straightforward. However, a new class of issues arise having to do with how the user's local view of structure is created and changes.

There are two main ways in which a user's local structure might arise. First, it might be explicitly constructed by the user, perhaps starting from one of a set of standard structures. This roughly corresponds to making a bookmark file a starting point in the world wide web today. Second, it might be based on the user's patterns of access in some frequency-based way. For example, 'trees' that were often visited might migrate into each user's local wood; and they might arrange themselves so that the most frequent trees were nearest to the notional initial viewpoint.

The issue of how a user's local structure changes depends, to some extent, on how it is constructed. The two alternatives for creating it can also be used to alter it; explicitly, by user action, or implicitly, based on usage. There is no need, even, for the same technique to be used for maintenance as was used for construction. However, there is a third possibility – forbid routine changes to the structure other than growth at the edges. The argument for this is that one of the reasons to use this style of conceptual model is that it builds on humans skill at remembering spatial navigation. Keeping the local structure relatively fixed allows spatial memory to play its role in *finding again*. Forbidding changes is probably too strong; but the essence is that making changes should be difficult enough to become memorable so that they stick in the user's memory.

In the global structure case, the methodology for defining the global structure defines how each local structure blends into the global structure. When each local structure is individual and specialised, this question becomes more difficult. Some of the obvious possibilities are: using organisational or technical structures, so that each user's local structure is adjacent to his or her natural neighbours; or using a loose global structure into which each user can slot as desired ("trails through the forest").

7 Assessing Treeworld

Is the meta-structure *visible* or *implicit*? Treeworld's main difference from other large-scale hypermedia conceptual models is that it provides an explicit meta-structure. This introduces the concepts of the inside and outside of a hypermedia document – from the outside, a rendering of the local neighbourhood is visible; from inside the content of the document is visible.

Does the meta-structure make *finding* easy? The advantage of Treeworld over other hypermedia systems is that new documents can be accessed by their proximity, position, and attributes in the rendering, without the need for any explicit linking or meta-linking via a search engine.

Does the meta-structure make *finding again* easy? Humans have good memories for spatial navigation. The multitude of subtle cues provided by shape, position, colour, and other attributes make it likely that previously visited nodes will be remembered.

Does each node's *context* in the meta-structure correspond to at least the context assumed by its author when it was constructed? The hierarchical structure of nested directories is one of the Unix ideas that has migrated widely into information systems (and hierarchy is, in any case, a natural human concept). Basing the rendering of a document's context on its position in such a hierarchy catches part of the context implied by the author. A wider context comes from the rendering of complete web sites, and the way in which each site is placed relative to other sites.

Does the meta-structure make *browsing* easy? Hypermedia systems without meta-structure limit browsing to whatever was conceived by the author. In Treeworld, the rendering of the local environment makes completely unrestricted browsing possible. If the ability to move through the environment in an unrestricted way is included, then the scope of browsing is unlimited.

Are the *actions of others* visible in the meta-structure? In existing hypermedia systems, each individual user is alone. The use of social cues indicating, for example, interest levels in a particular document are not naturally representable. A meta-structured system has a wide range of options in presenting the actions of others, from explicit presence to abstractions. Balancing quality of representation and computational cost are clearly important, and privacy may also be a concern, but Treeworld permits the full range of options.

Does the meta-structure use 'screen' real estate well? Limited screen real estate is best handled by some kind of focus+context approach (fish eye, fractal) since it provides both directly accessible information and a sense of what wider information is available. Trees (and woods) are an example of this kind of representation because they are scale independent. Apart from being recognisable at different scales, their shapes can also be used to convey information.

Does the meta-structure use a *metaphor* that is appropriate and natural for humans? Spatial metaphors play to human perceptual strengths, in recognition and memory, and are direct representations of the real world and hence natural arenas for action. Treeworld uses the metaphor of a movement in a wood, an almost universal human experience.

Where does the meta-structure come from? There is increasing pressure for hypermedia documents to announce their type to provide a foundation for inferring meta-structure. At present, type information is almost completely absent. In Treeworld, we have suggested that the structural information implicit in the url of each document serve directly to provide the meta-structure. This is rather weak, and perhaps even misleading, but it is a reasonable compromise. As meta-information about individual documents improves, more sophisticated schemes can be used.

Does the meta-structure *change* and, if so, over what time frame? The need to find again should exert a considerable brake on changes to meta-structures. In Treeworld, we anticipate that changes will take place on relatively long time scales: perhaps days for individual nodes in a tree, and months for changes in the position of tree themselves.

Does the meta-structure make *effective use* of size, position, colour, and movement in rendering its information? Since little rendering of web meta-structures has been done, almost nothing is known about this area. Experiments with VRML have been disappointing both from the point of view of the quality of the rendering and the cost of the computations needed to generate it.

8 Conclusion

We have presented the Treeworld conceptual model for complex, large-scale hypermedia that merges developments in technology, such as cheap 3-d visualisation devices, with the way humans move around the real world to produce a visual interface that is more flexible than current browsers but more lightweight than virtual reality. Treeworld is a framework for exploring the impact of design decisions about meta-structure, navigational metaphors, and visual interfaces.

The computational needs of the model and rendering needs of the interface outweigh the low cost of the interface hardware at present, but computation has been getting cheaper for a long time. The approach may well become practical within a few years.

Acknowledgements

David Skillicorn is supported by the Natural Sciences and Engineering Council of Canada. Kevin Brewer implemented a prototype Treeworld rendering system.

References

[1] K. Andrews, J. Wolte, and M. Pichler. Information Pyramids: A new approach to visualising large hierarchies. In *IEEE Visualization'97*, pages 49–52, October 1997.

[2] L. Beaudoin, M.-A. Parent, and L.C. Vroomen. Cheops: A compact explorer for complex hierarchies. In *IEEE Visualization'96*, pages 87–92, October 1996.

[3] M. Berry, S.T. Dumais, and G.W. O'Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37(4):573–595, 1995.

[4] M.H. Brown and R.A. Shillner. Deckscape: An experimental web browser. Technical Report 135a, Digital SRC, March 1995.

[5] M.H. Brown and R.A. Shillner. A new paradigm for browsing the web. In *Proceedings of CHI'95*, 1995.

[6] A. Dieberger. Browsing the www by interacting with a textual virtual environment – a framework for experimenting with navigational metaphors. In *Proceedings of Hypertext'96*, pages 170–179, 1996.

[7] J.D. Donath. *Inhabiting the Virtual City: The Design of Social Environments for Electronic Communities*. PhD thesis, Massachussetts Institute of Technology, 1996.

[8] B.A. Huberman, P.L.T. Pirolli, J.E. Pitkow, and R.J. Lukose. Strong regularities in world wide web surfing. *Science*, 280:95–97, 3 April, 1998.

[9] B. Johnson and B. Shneiderman. Treemaps: A space filling approach to the visualization of hierarchical information. In *Proceedings of IEEE Visualization'91 Conference*, pages 284–291, 1991.

[10] Frank Kappe, Keith Andrews, Jörg Faschingbauer, Mansuet Gaisbauer, Hermann Maurer, Michael Pichler, and Jürgen Schipflinger. Hyper-G: A new tool for distributed hypermedia. In *International Conference on Distributed Multimedia Systems and Applications, Honolulu, Hawaii*, pages 209 – 214. IASTED/ISMM, August 1994.

[11] J. Lamping, R. Rao, and P. Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of SIGCHI'95*, pages 401–408, 1995.

[12] T. Matlock and D. Modjeska. An experiential approach to getting around in virtual information spaces. In *Proceedings of 2nd Swedish Symposium on Multimodal Communication*, October 1998.

[13] S. Mukherjea and J.D. Foley. Visualizing the world-wide web with the navigational view builder. *Computer Networks and ISDN Systems, Special Issue on Third International World Wide Wen Conference*, 27(6):1075–1087, April 1995.

[14] S. Mukherjea, J.D. Foley, and S. Hudson. Visualizing complex hypermedia networks through multiple hierarchical views. In *Proceedings of SIGCHI'95*, pages 331–337, May 1995.

[15] R. Nideffer. Public spaces and meaningful traces: Interfacing to information personae. In Thomas Swiss, editor, *Magic, Metaphor and Power: The WWW and Contemporary Cultural Theory*. to appear.

[16] P. Pirolli, J. Pitkow, and R. Rao. Silk from a sow's ear: Extracting usable structures from the web. In *Proceedings of SIGCHI'96*, 1996.

[17] G.G. Robertson, J.D. Mackinlay, and S. Card. Cone Trees: Animated 3d visualization of hierarchical information. In *Proceedings of SIGCHI'91 Conference on Human Factors in Computer Systems*, pages 189–194, 1991.

[18] E.D. Smilowitz. Do metaphors make web browsers easier to use? `www.baddesigns.com/mswebcnf.htm`.

[19] L. Tauscher and S. Greenberg. How people revisit web pages: Empirical findings and implications for the design of history systems. *International Journal of Human Computer Studies*, 47(1):97–128, 1997.