Navigation and Orientation in Hypermedia

The Geographical Metaphor

People use metaphors in everyday life to make it easier to understand little-known concepts. They are an integral part of everyday thought. People try to understand unknown concepts through comparing them to something known that might exhibit similar characteristics or behaviour. "*Metaphors function as natural models, allowing us to take our knowledge of familiar, concrete objects and experiences and use it to give structure to more abstract concepts*." [Erickson, 1990]

Various metaphors have been used on the computer. A successful example is the desktop metaphor used on the Macintosh computer. It attempts to hide the complexity of the operating system by presenting a consistent, familiar image to us. It represents our electronic documents and applications (or tools) as icons that appear similar to small pieces of paper. These documents and tools are available on a desktop. Tools and documents can be stored in folders on the desktop, and they may be moved around by directly manipulating them. A document can be thrown away by placing it in the trash can. Of course, all this is a metaphor — the electronic documents do not actually move, pointers to them are moved, but this would be too complex and difficult to understand for most users, so the interface metaphor works to a large extent.

Some exceptions continue to cause problems, such as the ejection of a disk by dragging the disk icon to the trash-can. This is where continuing the metaphor has problems — what can a disk be associated with on a real desktop? Perhaps the disk should be ejected by placing it in a filing cabinet or disk box? This would be continuing the desktop metaphor if the disk was thought of as a folder to be filed away, whereas putting it in the trash would seem to be deleting what was contained within it. The success of a metaphor is attributable largely to the correspondence in appearance, use and behaviour of the interface 'objects' — documents and folders — and their real-world counterparts. Often users' problems are in the differences in behaviour, or where the 'mapping' between objects or situations breaks down. Successful metaphors should emphasise certain features and suppress others [Mountford, 1990].

Information Space and the Geographical Metaphor

The quantity of information that is available is often likened to the 'world of information', 'information space', or the 'information landscape' [Florin, 1990]. This world of information has often presented problems for people who want to find information within it. The previous physical nature of much information — that printed in books — and its subsequent transferral to electronic media has resulted in the actual geographical problems associated with finding information in the physical world being compared, through metaphor, to the problems of finding information using the computer in the electronic 'world'. The nature of finding one's way in the real world is also likened to that of finding information, because a place in the world *is* information — it is actual physical substance and not an abstraction through words or symbols.

The appearance of hypermedia systems has brought increased usage of the geographical metaphor for accessing information. The concepts of nodes and links have resulted in metaphors such as *following* a link, and when a number of links have been followed, a *path* has been followed. *Landmark* nodes are often created that can be rapidly visually distinguished from the surrounding nodes or space. This is obviously using a concept that is familiar to almost everyone — the idea of being or going somewhere, and this in turn uses the idea of landscape or geography.

> *"The terrain on which the information landscape is built is the raw database, rich with various materials ... the information structure is what gives the landscape its distinctive features."*
> [Florin, 1990, p. 31]

Parunak (1989) describes five different topologies that might be used in the hypermedia world and describes the navigational strategies that would apply to each topology. From simple to complex, these topologies are: linear, hierarchy, hypercube, directed acyclic graph, and arbitrary. Different users apply different strategies when navigating these topologies. The main navigational strategies used in the real world, and which may be applied to other complex systems such as hypermedia systems, are:

— Identifier Strategy — associates a unique identifier with each entity of interest thus permitting the searcher to recognise the target.

— Path Strategy — provides a procedural description of how to get to the target.

— Direction Strategy — uses a global framework as well as the ideas of texture and comparability. ‘Texture’ is the existence of a reference point relative to which directions can be established. ‘Comparability’ is the existence of a relation for any two points in the space.

— Distance Strategy — bounds search to a circle (based in time or distance) around the traveller’s current location. Often used in conjunction with the direction strategy, but becomes degenerate when any point is accessible directly from any other point.

— Address Strategy — refines the direction strategy by establishing an orthogonal set of coordinates, such as longitude and latitude, or the grid formed by streets in a town.

Of course, these strategies are often employed together. The number of these strategies that can be used reduces as the complexity of the topology increases. So when there are complex topologies as in many hypermedia systems, ways must be found to reduce the complexity so that more strategies may be utilised in order to effectively navigate to items of interest.

One way to do this is to impose some other structure on top of the system. Another way, which is complementary to additional structure and continues the geographical metaphor, is to use tools that are familiar to us from real-world navigation, such as a map or a beaten-path mechanism. A beaten-path mechanism might just be a history function that remembers the nodes that have been visited and allows the user to backtrack easily. A map can be a two-dimensional representation of the information landscape and can be used at a number of levels as in physical geography.

Other geographically based metaphors have been used in some circumstances. Hammond and Allinson (1987) used a tourist and travel holiday metaphor as a design aid in helping users to navigate complex systems. They have extended the idea to learners getting on a ‘tour bus’ to follow a guided tour through the system [Hammond and Allinson, 1988]. The tourist and guide metaphors have also been used by Fairchild et al. (1989). The systems that these have been incorporated into have all been relatively small. A large scale system that utilises the geographical metaphor is the Hyper-G system which extends these ideas into the ‘world of travel’ metaphor [Davies et al., 1991]. This enables them to bring many different metaphors — such as maps, guided tours, paths, agents or guides for varying purposes (e.g. travel agent, tour guide) — under one unifying theme. This can help users to relate what they know about navigating the real world to navigating the world of information. When they need help finding something they might go to a travel agent who can suggest where to go to find it, or the agent might suggest a guided tour of an area if the user is a novice. There might be information guides, similar to travel guides, which give advice about what to see, where to find it, and what value it has — this is like getting an expert’s opinion on an area and it's probably more valuable than asking your next-door neighbour.

The path facility described in this thesis continues the geographical metaphor. A guided path is used as a mechanism with which anyone may tour through information in a structured way (i.e. following the path). Each node can be considered a landmark, and each link can be considered the path between landmarks. At any landmark side trips available from it can be taken, and then the landmark can be easily returned to when the side trip is finished so the tour can continue. Thus the structural and conceptual simplicity of a linear path can be combined with the possibility of structurally complex side trips at any point.

Navigating Hypermedia — Overview

How is information accessed in a hypermedia system? There are many methods and tools that are employed which can be classified into three main groups: large scale, medium scale and small scale. The large scale tools are used with large-scale hypertexts — those which encompass multiple sites or ones that span the globe, for instance. Medium scale tools are used with medium scale hypertexts — those based at a single site, for example. Small scale tools are those which are used to navigate a single hypertext document.

Large Scale Hypertexts

In the past, most effort has been aimed at developing tools for navigating small and medium scale hypertexts, mainly because those are the main kinds that have been developed so far. There are very few, if any, large scale hypertexts as yet, though the problems inherent in large scale hypertexts are being considered. The most ambitious large scale project is Ted Nelson's Xanadu project which aims to be "*servicing hundreds of millions of simultaneous users with hypertext, graphics, audio, movies, and hypermedia*" [Nelson, 1988].

Another large scale project is the WorldWideWeb (WWW) project being developed in Switzerland. Its aim is to increase the accessibility of academic information, "*... to allow information sharing within internationally dispersed teams, and the dissemination of information by support groups*." [Berners-Lee, 1991] It uses tools such as indexes, search facilities, and browsers. All documents look the same to the reader and may be accessed in two ways — through an embedded link, or via a search mechanism. It is stated that "*These are the only operations necessary to access the entire world of data*." [Berners-Lee, 1991]. That may be true, but in order that the system be readily usable, other orientation and navigational support tools must be built into it. Merely having access to data is quite different from being able to access it easily. It also ignores the need for visual cues that are useful in the various levels of navigation.

A very recent project is the Wide Area Information Servers (WAIS) project [Kahle, 1991]. This uses both static and dynamic linking to access information from all over the world using a standard protocol. Navigation is achieved through a search mechanism using a natural-language type query. The words in the query are matched to documents, with articles that have the most words matching being ranked highest in importance. There are a variety of other navigation mechanisms being a variation of this main theme.

Other navigational tools could be provided so that these systems are more easily usable. For example, a map or globe facility would be helpful in determining a number of things about the information, such as cost, and time to access it. Details such as these are important for navigation in the large because if we are going somewhere to get something, it should be known how long it will take and how much it's going to cost.

Medium Scale Hypertexts

Medium scale hypertexts offer a wide range of navigation and orientation tools, and these tools vary from system to system in their implementation. Almost all, however, offer the two main types of tools — a graphical browser, and a search facility. These vary widely in their implementation, however, and there are also a number of other tools used for navigation in medium scale hypertexts. These tools will be illustrated from two of the most well known current hypertext systems — Intermedia and NoteCards. Subsequently a few other tools or tool refinements that could be useful for these systems will be suggested.

Most current hypertext systems offer some type of graphical browser with which to navigate the system. It is somewhat analogous to a paper map although on the computer more flexibility exists than on paper so most browsers are automatically generated to give a sense of context. In some manner they attempt to show current location, past locations, and possible future locations. In the past, some browsers attempted to show all nodes and links in the hypertext, but this was soon found to be unrealistic because as the number of nodes increases, the number of links increases exponentially [Baird, 1988]. This quickly makes for spaghetti-like maps, so some form of selection and filtering must take place in order that a browser be usable by an individual.

Another problem with browsers is that they seem to be needed on a number of levels. First there is the local browser which shows your immediate context within some document. Then there is an overview browser which might attempt to show medium-scale context — that is, it might show the nodes at the current site and some linkages between them. Then perhaps there will be a need for a global map, which shows how systems around the globe are linked.

The first two browsers could perhaps be implemented in one using some sort of ‘fish-eye’ browser [Furnas, 1986], which shows a high amount of detail for nodes in the immediate vicinity, but the detail becomes increasingly less as the distance from the point of interest increases. So a sense of our immediate context can be gained, but also some idea of context in a larger sense. Local detail can be seen in a global context.

The global browser may be more difficult to implement but might be helpful in a number of ways. To begin with, it would serve to give an indication where repositories of certain types of information were located. It could also serve such purposes as ‘information demographics’ much as a geographical map displays geographical information. Perhaps the map could change to display relative densities of various types of information. Hence if a particular area was interested in, a filter could be applied to display where the main centres for this information were using the global map. Then perhaps a particular area could be zoomed in on and the overview browser or map for that area would appear. This might enable us to quickly locate areas of relevant information. It is a method of filtering the information displayed so that it is relevant to us. Thus as the filter is defined more precisely, the displayed map becomes even more meaningful. The ‘zooming in’ to display different levels of a map or browser illustrates the multiple levels needed in order to get immediate, local and global context. There can also be different levels within these arbitrary divisions.

The information shown in these graphical browsers can vary. Some are text-based with each node being represented by a short description of its contents. Some are graphically based with a node being represented by an icon. Generally these icons represent the type of node, for many systems allow typing of nodes, e.g. a node might be text, graphic, sound, video, animation, or some other type, and each will have its own iconic form on the map. Often there will be a combination of the two with an icon representing the type of node and a short textual description of it. Of course, in a hypermedia system, each node should probably have a short representation of its contents. For example, a graphic might have a thumbnail of the graphic, or a node containing video might display a short small video clip. This would conform to the user-interface principle of progressive disclosure, although in a system with an unlimited ‘go-back’ facility and quick response time this might seem less necessary.

Intermedia offers an interesting and valuable tool called the web [Utting, 1989]. This is a device that helps us get context by filtering out unnecessary information. A web can be opened and within it only those links and nodes that belong to that web will be available. Therefore it simplifies the graphical browser, especially where multiple links emanate from nodes, as it can only show the nodes and links that are applicable to this web.

It might also be possible for us to create our own webs so that when the web is returned to at a later date, a similar sense of context can be gained through following the same path that was taken previously. This would seem to be quite desirable, for although the seemingly unstructured nature of hypertext is deemed as one of its major strengths, when an argument is developed a line of thought is followed. So a web that replays a linear path through the structure can help as a reminder of what was previously gained from it. Or it could be used to apply linearity for someone else to read. Then it is a form of guidance which may be used to help a novice through the difficulties involved in navigating complex structures.

The other main medium-scale navigational aid is a search mechanism, which enables users to break through structural boundaries to locate information. A search is often used when a particular item, which satisfies some criteria, is being looked for and its location is unknown. So a rapid scan of the system may be used to check for the existence of the required items. Often an index will be used within the search mechanism. A search results in a number of node hits which can then be evaluated as to relevance, or the nodes can be immediately accessed. A hit can be classified as an occurrence of the item that satisfies the search criteria.

In the Dynamic Medical Handbook Project [Frisse, 1988] a node's relative value is calculated from two components. The intrinsic component is calculated from the number of hits within that node, while the extrinsic component is a value computed from the weight of the node's immediate descendants. The sum of the two components gives a node's value that can then be displayed, perhaps on the browser, so the nodes which are the most relevant to our query can be immediately identified. In this case the query could serve as a filter and display only those nodes that are hit. When these individual nodes are accessed, any descendent nodes from there may also be accessed, not only nodes that contain hits.

Other search mechanisms use various weighting methods to assess the relevance of a particular document or node to the search query. The WAIS system uses simple text pattern-matching algorithms so that the number of hits of the same text within a document will be reflected in its ranking [Kahle, 1989]. Other systems use sophisticated searching techniques that can also base searches on the structure of the system. For example, when a system uses typed links, as in NoteCards, the user might like to find nodes that both contain some particular text and are connected to other nodes by a particular type of link [Halasz, 1988]. This is a very powerful search mechanism, but one that requires much knowledge about the system by the user so it is more relevant to specialised systems. This does not mean that it shouldn't be provided in all systems of course. (For an extended description of search techniques, see [Salton, 1989] or [Ellis, 1990].)

Another important tool used in medium scale hypertexts is the Landmark. It provides navigational information through being a recognisable reference point from which a sense of orientation can be gained. That is, it can be easily returned to and, from it, other previously visited nodes can be returned to. This illustrates a need for having regular landmarks throughout a hypertext, and there are a variety of ways to make a node a landmark.

One way of doing this is to use different coloured or textured backgrounds for an occasional node. Of course, for this to be effective it should not be random. Rather it should be a node whose contents are somewhat different from its surrounding nodes. It could, for example, be a section heading. Nielsen [1990a] uses different textured backgrounds and different graphical designs in his hypertext to distinguish a regular node from other elements of the system.

Another way of providing a landmark in medium scale hypertext is to use reference nodes. These might be nodes which are in some way relevant to the subsequent nodes, a section heading or table of contents for example. They might be nodes at an upper level of a hierarchical structure. Often part of the reference node might be carried on to subsequent child nodes — subsequent nodes of the same section — to indicate the underlying structure. This might be the carrying over of a section heading, with it being in a consistent location throughout the section. This would assist the user in orientation and thus reduce the cognitive overhead involved so the user can concentrate on the content and not on working out the structure or what the relationships are.

A reference node should be a node that can be returned to very easily and which is immediately recognisable through some distinguishing characteristics. These distinguishing features might be visual cues such as colour, spatial layout, and background pattern, or perhaps it could be a feature such as an audio sequence, or animation, or more likely a combination of these. A combination of cues is needed because as a system gets larger, multiple cues are needed to distinguish one landmark from the next [Nielsen, 1990a].

A table of contents seems to be another valuable tool for medium-scale hypertexts, because it presents an introduction to the system contents so that the user can quickly determine whether this area is relevant. It can also aid in conceptual understanding and the construction of mental maps of the system structure [Simpson and McKnight, 1989] which can increase navigational efficiency.

An index is also a valuable aid as it can provide direct access to an item that is known to be in the system, when it is not known exactly where in the system it is. This form of index is similar to an index in a book, which contains items that are deemed to be of interest by the author and pointers to their locations in the document. It differs from the index associated with a search mechanism in that a search index will often contain all information in the system in an indexed form, and this will not be in a form that can be directly used. It must be accessed through the search interface which will allow the user to construct a query. The book-like interface will contain the author-specified terms of importance and their locations, and these terms may then be directly accessed within the system.

Small-Scale Hypertexts

Navigating hypertext in the small uses a variety of tools and techniques, most of which are visual cues that aid our understanding of the hypertext structure, node content, and screen layout. Thus they are aids for understanding the current position in relation to the medium-scale (i.e. the current structure), how to find information on the current screen, and how the current screen's information is structurally arranged. Standards also have an important part to play in small-scale hypertext navigation because, once the standards are learnt, they allow the user to concentrate on the information presented rather than dividing the user's attention between presentation and content.

Thus small scale tools are primarily visual cues and conventions, much as standardised conventions such as page numbering and paragraph arrangements exist in printed media. Kahn et al. (1990) identify three graphic design principles that are appropriate for the design of hypermedia documents and relevant to small-scale navigation:

The rules of type — the relationship of type, leading, and line length to legibility. To maintain legibility, adjustments must be made to account for the low resolution of the computer screen.

Consistent formatting — single publications or series of publications should contain consistent formatting rules to support reader orientation.

Clear information graphics — again the low resolution of the computer screen should be taken into consideration.

However, a hypertext offers not only more flexibility than normal documents but also potentially more problems. With the flexibility of presentation come additional problems in understanding what is presented. One particular area of interest is the question of whether to present many windows per screen, or only one window per screen. An extension of this question is whether to present one idea on a scrollable window, or whether the window size should be fixed and if the idea is too large for one window, then it should be split across multiple windows.

If one node is presented as a single idea, how should this be implemented? One of the main aims should be to make it clear how much information is contained in the node so how much more there is to read, look at, or understand is easily known. The implementation of this might vary. One very common way is to present a paragraph on a single screen, with many screens making up one idea or concept. Another way would be to have a scrollable window presenting one concept on one screen.

The main issue involved here is the presence of cues that enable users to ascertain how much information there is, how long it might take to go through it; how a previous position can be returned to; and how to find something of interest within the idea presented.

In the many screens per idea approach, page numbering schemes can be used as a cue, with an example being the words ‘Page 1 of 6’ at some position on the screen. Or a visual cue such as that used by Nielsen in his Hypertext ‘87 Trip Report [Nielsen, 1987] could be used. He uses a version of a horizontal scroll bar to indicate the relative position through the current idea. A type of horizontal scroll bar can be used to move forwards or backwards to a position, allowing rapid re-positioning through the hypertext. It is also possible to move linearly through this section of the hypertext through the next page and previous page icons.

In the one-screen-per-idea approach a scrolling mechanism might be used with other visual cues to indicate position. One way would be to use modified scroll bars such as used in the Atari computer rather than those used on the Macintosh. These indicate the amount of information being currently presented relative to the whole so that how much information there is, where the current position is within it, and how much more there is to go within this node, is easily ascertained.

One possible problem with the one-screen-per-idea approach is that the micro-cues that are obtained from a single page may be more difficult to understand. For example, if it is desired to quickly browse through a section searching for a previously known location, then cues such as page layout and relative white space (the layout of paragraphs and sentences), that can distinguish one page from the next, have probably been lost. With the many pages per idea approach more cues are available so it becomes easier to navigate. This assumes of course that the cues presented are unambiguous.

Summary

The flexibility offered in accessing information through hypertext systems must be accompanied by various tools that enable all levels of users to access the information efficiently and effectively. There are a number of levels of navigational tools and cues that must be provided for such access to occur, and these can be categorised as large-, medium-, and small-scale. Some concepts are present at all levels, such as landmarks and queries, while others are particularly suited to one level or another, such as a table of contents which is used for medium-scale hypertexts. Many current systems provide medium-scale tools because the systems are designed for that level. The small-scale presentational conventions are usually left up to the individual designer and this is somewhat unfortunate as a multitude of styles leads to incompatibility and difficulty in understanding. The systems that are designed for the large-scale usually neglect the medium- and small-scale navigational aids, and these need to be provided for these systems so that accessibility and ease of use are increased. An understanding of the multiple levels of navigational mechanisms is required to provide effective usability for the differing requirements of many users with differing skills and experience, and to provide efficient access through the different levels or ‘scales’ of hypertext systems.

The Path facility will be an effective tool in aiding navigation for some users at all levels. It can assist users to gain overviews of different areas, to navigate large systems easily, to store paths within large systems and so efficiently navigate back to previously visited areas. It will help a user with small-scale navigation through the provision of meta-information that can alleviate some problems by providing specific directions on what to do at a node, and also an explanation of the conventions used at a node.

Problems

As with any system, casual users will often have problems with basic tasks. In a hypertext system, the main problem that users have is one of navigation. This manifests itself in questions such as:

Where am I?

What can I do here?

How did I get here?

Where can I go, and how do I get there? [Nievergelt & Weydert, 1980].

Is there anything about 'X' here?

I know there's something about 'X'— how do I find it again?

How much is there here?

These problems are to do with navigation and orientation, and can result from a variety of deficiencies in the system itself or in the design of the hypertext documents.

It is very difficult to structure a hypertext system so that it can satisfy all users and all tasks. Consequently it is often difficult to provide an appropriate overview of the material — some users may fail to see the structure of the information base, and so may miss out complete sections. The number of potential links and paths is sometimes great, so a user may become overwhelmed with choice and consequently freeze up. Also, link markers may not always be obvious, and this can result in the user missing potentially valuable information.

Providing a graphical browser seems to be important in assisting users to gain some sense of their context, or their surrounding information space. Edwards and Hardman (1989) conclude that a 2- or 3-dimensional representation of the information structure is an appropriate navigational device. However, problems exist with such displays when the number of nodes and links increases and the structure becomes very complex. This illustrates the need for other tools and devices, both to assist in filtering the structure and to provide additional context.

Link Markers

An initial problem that many users have is in discovering where links exist. How are the links in a particular system distinguished to the users? And what is the extent of the link — that is, what part of the node is the link origin? For example, if some text in a node is linked to another node, how is it indicated to the user that a link exists, and secondly, how is it shown what part of the text is linked?

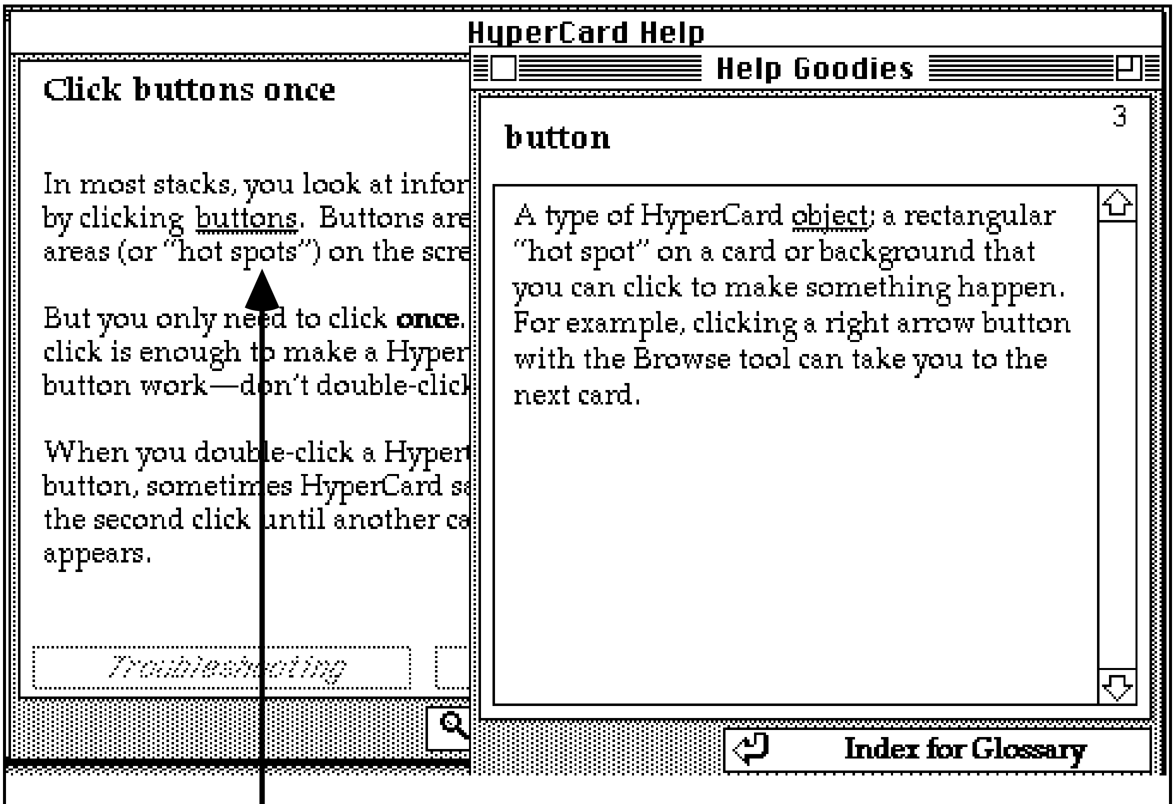

Note the greyed underline of the word 'buttons'. This indicates that the text is in the style 'grouped' and is linked. This picture show the result of following the link — the destination node is shown at the right in the window 'Help Goodies' and is the card with 'button' information on it.

Figure 2.1 *Link anchor, anchor extent, and destination using HyperCard grouped text as an example. The grayed underline of grouped text can indicate both the existence of a link and the link extent at the same time. This is useful if this convention is immediately recognised.*

In some systems there are standard conventions for this. Often some form of highlighting of the linked item occurs. A textual link might be in a different type style or a different colour for example. This is an easy solution that quickly indicates the link existence as well as its origin. But this might conflict with other forms of highlighting such as that used to add emphasis. Evenson and

Rheinfrank (1989) suggest that typographers might design an entirely new set of characters that would signal ‘hyperness’, but which wouldn’t detract from readability. Unfortunately this does not encompass other media such as graphics. Typographical styling is no help there, so another solution is needed.

Evenson and Rheinfrank go on to suggest, however, that some sort of new hypermedia design language using an action-based sign-set would complement the existing alphabetic symbols that are currently used. The new sign-set could be provided with guidelines for combining them in order to create visual cues that could result in showing not only that a link exists, but also the action that will result on following the link. That is, the link marker indicates the link origin and its type. The extent of the link might be indicated by an outline appearing when the link is selected. This has potential problems such as the initial extra overheads in understanding the types of links, but it does seem to be a promising solution for link marking as well as providing extra meanings to the visual cues.

Link Destinations

Just as in specifying what is the link origin, how is the scope of the link destination known? Does it relate to the whole destination node, or to just a part of it? This is a problem of the granularity of the system — how large are the ‘chunks’ of information that are connected? The smaller the chunks can be, then the more meaning can be taken from the linkages. If only links from one node to another are available, then the nodes must be made very small so that confusion as to the reason for linking is avoided. One way to do this is to have a multi-window system such as NoteCards where very small nodes are common. Another method is to use larger nodes but with the ability to link up smaller parts of the nodes with small parts of other nodes.

Backtracking and History

Another problem is how to return to a previously visited node. Both of these problems may be partially solved through the use of a rapid return or go-back facility. This is really a prerequisite in any computer system that encourages exploration — something may be attempted and if it’s not what is desired, then it is easy to return to the previous state. In this case, a link may be followed and if it is decided that it’s not relevant at this time, then it is possible to go back to the previous node. It provides a safety net that encourages exploratory learning [Carroll, 1982].

A supplement is some form of link-previewing or progressive disclosure, where some information about the destination will be provided before actually following a link. This may take many forms depending on the type of the destination node. For example, if the node were a graphic, then perhaps a reduced view of the graphic could be shown. If the destination were a text node, then perhaps the title and length of the node could be shown. If the node were an animation or video then a reduced version of it could be played, reduced both in time and in size. These forms of link previewing are beginning to be shown in some computer systems such as the Macintosh with the QuickTime extensions to its system software. It is only a matter of time before they begin to appear in hypertext systems.

Scope

Another problem associated with navigating systems is the problem of scope, or how much information there is. Should there be an indication of how many nodes and links are in this system, this web, or this document? Which navigational strategies should be employed if the extent or size of the system is unknown? If you are in the middle of following a path, for example, and you want to finish soon, how do you know how far there is to go to the end? Should you stop now and come back later, or are you near the end already? This indicates the need for some global status information.

Meta-information

The guided tour idea has also presented problems in navigation and in understanding. Because a guided tour is system-controlled, a user has, to some extent, ‘lost’ control of where they may go. This means that a user will often want to know why they have been led to a particular node. It seems that meta-information — information about the structure of the tour — is needed to help make the tour intelligible and to help the reader avoid disorientation.

Meta-information needs to be distinguished in some way to differentiate it from normal content. The use of different fonts, size and style, and spatial location can be used to indicate this.

Summary

In summary, problems involved in navigating hypertext result in user disorientation and manifest themselves in such questions as:

Where am I? Where can I go from here? What can I do here? Is there anything about <topic> here? How much is there about <topic> here? How do I get to <topic>?

These questions arise from the complexity of the structure of the system. In small systems, the problems can often be easily dealt with through the imposition of some structure, but in larger systems when the number of nodes and links is great, structural imposition is one part of a larger number of solutions. One solution is to provide tools that assist the user to gain a sense of the structure. These might include devices such as graphical browsers, maps and landmarks, as well as page-layout conventions such as headings and scope information. Other aids include:

— devices that can simplify the structure of the system such as paths, tours and filters;

— meta-information that provides explanations of conventions used in a node, as well as context information through explanations and points of view;

— backtracking facilities that provide a safety net for users who explore the system;

— search facilities that enable a user to bypass the structures of the system to locate information.