Chapter 2

Background

2.1 Preamble

Information spaces such as the World Wide Web are growing larger, becoming more diverse, and more dynamic. Users increasingly need assistance to locate the information they require. In an information space that promises, but does not guarantee, that the information users are looking actually exists, users either need to be guided to the information, assisted in searching for it, or told that it does not, in fact, exist. Failure to provide this level of assistance could lead to users not locating information that does exist, or else spending inordinate amounts of time determining that it does not exist, '...like the blind man, in the dark room, looking for the black cat, that isn't there' [24].

As users become familiar with the portions of an information space that they frequent (particular Web sites, for instance), they become confident, and are usually able to rapidly locate the information they require, through either browsing or searching. Among the reasons for this gain in efficiency are familiarity with the layout, or organisation, of information carried by the Web site (which aids browsing) and the terms of reference used within the site (which aids searching). Parallels may be drawn with our everyday interaction with real-world environments. Consider the problems faced by a newcomer to a town or small city.

When a person moves to a new town, in order to effectively locate businesses, she must learn how the town is organised. If the newcomer has pre-knowledge that a newsagent sells newspapers, a supermarket sells food, a sports emporium sells sporting equipment, and so on, then the newcomer has simply to locate these outlets in order to make the necessary purchases and survive. There are several ways in which the locations of outlets can be found. The newcomer can buy a map; alternatively, she may decide to locate the

town's main street and simply traverse it - using the knowledge that each town has a main street which contains diverse businesses; or, the newcomer may refer to the town's business directory. When dealing with commercial entities in "real" towns, most cultures operate in a similar fashion, and, armed with the requisite fore-knowledge, any person can successfully locate commercial outlets within a short space of time. Locating other services, such as tourist attractions and plumbers, may take more time - maps and business directories are a source of information for the newcomer, but where can these be obtained from, if the locations of a newsagents, library and local telecommunications office are unknown? A solution could be that the newcomer hops into her car, or catches a taxi or bus, and drives into the town centre, the directions to which are usually well sign-posted. From there, an excursion along the main street will usually yield the information required to travel further afield.

This ability to navigate around town and locate the required services is learnt as one grows up in a similar environment. Signposts are located, and streets and businesses are named, to assist the navigator. If a street is not called "Main Street", then the presence of a large number of commercial outlets will indicate that the street is the commercial hub of the town. However, consider now the following scenario. Imagine that a town has no signposts, and that all premises, commercial or otherwise, are unnamed and have no distinguishing features, and, moreover, the occupants of each of the premises do not know from where services other than those offered by themselves may be available. The task of the information forager has become decidedly harder. Some clues about the location of the town centre may be forthcoming (such as the direction of flow of heavy traffic), but the newcomer is typically reduced to entering each building and asking either what the establishment is, or whether the establishment is able to provide a particular service.

2.2 Searching and navigating through information spaces

An information space without aids such as indices, maps, or distinguishing features presents a mode of interaction as unfriendly and unhelpful as the featureless (and, thankfully, imaginary) town described above. With small towns and small information spaces which contain features or for which we have a map, we can quickly learn what commodities (information or otherwise) are available and where the commodity sources are located. We require assistance only with the initial learning phase, and with being informed of changes to the community. Parallels may also be drawn between an information space which is large and contains multiple sources of similar information and a large town or city which contains multiple sources of commodities (for example, multiple newsagents, banks, clothes shops, and supermarkets). When a large town- or

city-dweller requires a service, there may be a number of criteria for selecting just which establishment will be visited to obtain the service. Amongst these are convenience, quality of service, customer care, value for money, and brand name or reputation. For example, the morning newspaper is likely to be bought from the nearest newsagent in walking distance, if it is to be read at home, otherwise it may be bought from a conveniently located newsagent on the way to work. If many items are to be purchased from different stores, it may be more efficient to visit a shopping mall or main street, where all the stores are located within walking distance of each other. If expensive items are to be purchased, the shopper may decide to visit a number of stores offering the same or similar items before buying the items which are within budget without compromising other requirements, such as quality. One supermarket may be preferred over another, not necessarily due to the brands carried by the supermarket in particular, but because shopping in the slightly more expensive supermarket is less stressful than battling it out with hundreds of shoppers in the discount store. Even though the selection process is more complicated in larger towns and cities, it is still reasonably manageable as each problem space is still relatively small. Whatever the task, the number of establishments offering that product or information will be small enough and probably within a relatively small geographical area for even exhaustive browsing to be tractable.

In large to massive modern digital information spaces, however, there is no concept of geographical distance - any information contained in the space is at the fingertips of the information seeker. Consequently, the number of sources of the required commodity is likely to be unmanageably large; so large that exhaustive browsing is not an option for the average information seeker. In towns and cities, commodity seekers find their way around using established signs which give information about their purpose to the seeker. The search is restricted to a reasonably-sized geographical region which can be quickly exhaustively searched, and the criteria on which decisions will be made are normally fairly rudimentary. We are considerably more demanding when searching for commodities in a digital information space. We are led to believe that the information we seek is actually out there - if we cannot find it easily, then the problem lies in a discrepancy between the description of the commodity we provide to a search engine and the description of that commodity that the search engine has, or because we are looking in the wrong place. We are also more critical of the information we find, because our expectations are different. We expect to find precisely the commodity we seek, rather than the best commodity from those available. We always have the feeling that there must be something more or something better out there, because our search space is not restricted, and because we have the expectation that the commodity does exist in the information space.

Some reasons, therefore, which make searching in and navigating through a digital information space difficult are: information providers have not yet learnt how to present and organise information so that it can be easily found; information spaces are not clearly delimited; we cannot always clearly describe the information we need; there is no guarantee that information we are looking for is described in the terms that we want to use; and, because of a combination of these problems, we don't know when it is safe to give up searching or navigating. These lead to associated problems of information overload and search overkill. Additionally, because modern information spaces tend to be large, dynamic, and fragmented, we have major problems understanding their organisation and we may regularly have to undergo the learning or re-learning process.

2.3 Organisation of information versus meta-information

Are things what they are because of their structural position in some closed universe, or for some other reason? And if so, what is that reason and can it be replicated in an information space? If we classify things then the two major tasks are to correctly place new things into some predetermined location and to guide a user to the correct sub-part of that space. On the other hand, if the information space is left unstructured, then some other information is required before users can be guided to the relevant information.

In a highly structured closed world, all information related to the same topic can be found in the same sub-structure. If a single piece of information is related to multiple topics, then it will simultaneously exist in different sub-structures. One can imagine hypertext links or references in such a connected environment seeming to appear and disappear depending on which of an information node's parent's was used for access. Consider a thesaurus which is converted into an structured information space which automatically partitions itself to reflect the word and word sense chosen by the user. Each node in the information space represents a particular word sense of a word and contains a list of words related to that word and word sense, with links to the nodes representing each related word in the list. Consider a user looking for meanings related to the word *iron*. Initially, as *iron* has several word senses, a list containing each of the possible word senses is presented, amongst them **noun** (as in the metal) and **verb** (as in to make smooth). Once the user commits to a particular word sense (follows the link to the appropriate node), links to words related to the other senses of the word "disappear". Similarly, if the user were to select *press*, one of the words which means *iron* in the sense *to smooth clothes*, then only alternative meanings for *press* which are consistent with that word sense and meaning would be presented. Thus, the words in the information space have been successfully partitioned. Whereas the choices in a non-partitioned thesaurus are potentially limitless, in a partitioned thesaurus the choices are

much smaller and the user cannot be distracted from locating meanings of words which are related, not only to the original word, but also to the word sense.

The effort in constructing a structured information space is in ensuring that the structure is correct and complete. Errors and omissions will, if they are significant enough, result in users either not locating what they need or being misled.

Consider an unstructured information space, again generated from a thesaurus, where each node contains a single word, its word sense, and a rule set for locating related words. Initially, when a user is searching for a word, all the word senses for the corresponding word are presented. Once the user selects a particular word sense, the rule set is triggered to locate related words. Here, the effort is concentrated in providing correct and complete rule sets. The rule set is responsible for partitioning the information space, which is constructed each time the user selects a word or word sense.

The advantage of the second approach over the first is that it is possible to add words and word senses to the thesaurus without knowing what other words already exist. In the structured information space, it is necessary to know how words and word senses should be connected to the existing information, whereas in the second approach, it is necessary to know only how to construct rule sets to locate related words, if and when those words exist within the information space. However, in order to be successful, standards are required to ensure that the rule sets are used appropriately and correctly. These rule sets form an intrinsic part of the information space, although they are not, strictly speaking, part of the information themselves. In other words, they are meta-information.

The World Wide Web (Web or WWW) follows the model of a structured information space, although it also has some characteristics of an unstructured space. Nodes are connected to each other via links, which represent relationships between information in the parent and child nodes. There are, however, two major problems with the Web. The first is that the actual semantics of a link is unclear. Links within the same node may represent a multitude of different meanings. Links may be navigational or else lead the user to related information. If the link leads to related information, which information in the parent is related to the information in the child is usually unclear. Is it simply the anchor text, the paragraph containing the anchor text, or does the scope extend beyond the paragraph? How much and what of the child document is relevant to the parent? The second major problem is that it can be virtually guaranteed that the parent node is not connected to all related information, because no individual knows all the information contained in the Web. Therefore, constructing the Web along the model of a structured

information space may be flawed, and consequently, users suffer problems with locating information.

Can the Web be remodelled so that information can be deposited into a largely unstructured space, with meta-information determining how nodes are related? When a user clicks on a link, rather than being taken to a pre-determined child node, a rule set is triggered to identify the most relevant documents, according to the criteria determined by the rule set. The rule set could also determine the scope of the link source and the type of relationship between the parent and linked documents. Of course, to be reliable, scoping and relationship standards are required, as well as an expressive enough ontology. Whereas new models of information space management may adopt this approach, it is unlikely that the Web will be remodelled (although recent proposals for the Semantic Web [5] may counter this statement!). However, there may be a happy medium between the two approaches to information management and user support, and this thesis considers such an approach.

2.4 Information re-use

We have suggested that a browsing space can be made smaller by using meta-information to partition the information space. Information which is not appropriate to a particular user can be hidden from view. It follows that if multiple users are browsing through a shared information space, they will each see a representation of the information space as it is relevant to themselves.

Consider a high-resolution 24-bit colour wall screen which carries information which is being constantly updated. The information may consist of the current time, weather, stock market information, and news headlines (Bloomberg TV [7]). Several users are simultaneously viewing the screen in order to glean information from it, although each has an interest in only some of the information. The information is carried in different parts of the screen. For example, the current time and weather is displayed in the top left hand corner of the screen. Although the time shown is Greenwich Mean Time, the weather information updates for the current conditions in major European cities. News is shown, one story at a time, in the top right hand corner of the screen. Stock market information scrolls along the bottom of the screen. People focus on the area of the screen which interests them. Apart from the current time, all other information is transitory, in the sense that items are replaced and shown again some time later (only a few stock prices are shown at a time, for instance). If the information of interest to a particular viewer is not currently displayed, she will have to wait until it scrolls back into view. The amount and type of information carried is limited to the amount of screen real estate available.

Consider a similar scenario, except that, although the screen is still shared, all information is carried simultaneously and is not replaced (so all stock information, weather conditions in all cities, and other information, are all shown permanently and concurrently), but viewers see only the information in which they are interested. All other information is filtered out of view. Viewers wear spectacles which respond to specific light frequencies and intensities. A viewer interested in a particular news item would wear spectacles which respond only to the light frequency and intensity used to display that news item. Each pixel on an RGB screen is represented by different intensities of the colours **R**ed, **G**reen, and **B**lue. Mixing these colours in different proportions yields millions of different colours or hues. If a particular news item is composed of pixel points each of which contain precisely 19% blue, regardless of the saturation of red and green, and the spectacles worn by individuals interested in that news item filter out all pixels which do not contain 19% blue, then that news item would be seen only by individuals wearing the appropriate spectacles , and those individuals would be able to see only that news item. Viewed without the spectacles , the screen would contain incomprehensible information. In this example, the meta-information is the strength of a particular R, G or B signal, to which the viewer's spectacles are attuned, which is not part of the information (such as a news item) itself, but without which no or limited sense can be made of the data.

Following this example, apart from partitioning an information space, meta-information allows information to be re-used. Each pixel can be used in at least three different patterns simultaneously - with different filters utilising the basic properties of the signal strengths in each pixel to determine solely whether a pixel will be made visible to the user.

A similar example, with which we are all familiar, is watching TV or listening to the radio. All TV and radio signals (within our catchment area) are picked up by the TV or radio aerial simultaneously. However, when we want to listen to or watch a particular station we simply tune our receiver into the appropriate frequency.

In an information space, the same information may also have different uses, or may be required by users for different purposes. This is made evident, in part, by the fact that the same document may have been reached by users navigating through different paths or submitting queries which contain different search terms. However, in the WWW and similar environments, although different users may have reached the same document by different routes, perhaps implying that the document supports different senses, the information space is not partitioned accordingly. In fact, regardless of how different users have arrived at a particular Web page, they will be presented with identical views of

the document and the structure within which it exists. Is this a disadvantage? Why might users have more difficulty locating relevant material than if the space was partitioned?

2.5 Navigational difficulties in an information space

Each time a user extends a path in a hypertext, she must make a decision about which link to follow, by selecting the link which appears to be most likely to lead to the information being sought. If documents are randomly linked together then any decision making process is futile, as a likely looking link is unlikely to lead to relevant material. The decision making process undertaken by a user is likely to be based on the assumption that the link's creator had a rational reason for creating the link, and that the material referenced by the link is somehow relevant to the link's description (such as the source anchor text). Two users who have accessed the same document, but who have followed different paths to reach it have in all probability reached different conclusions about why the document has appeared in their quests to locate relevant information. The document they have both reached contains the same set of links which lead on to other documents. Both users must again decide which link each will follow next. There are a number of possibilities. Both may decide that the document is in fact irrelevant and none of the links are likely to lead to relevant material. Otherwise, both may decide to follow the same link, although it is likely that the reasons for choosing the same link are different. Alternatively, they may decide to follow different links.

Consider two shoppers, armed with street maps, plotting routes through a town in order to each reach a different shop each. Their respective points of origin and their destinations are marked on their maps. Street corners are analogous to documents, and streets are analogous to hypertext links. Each time a shopper reaches a corner (a document) she has to decide which street (link) to follow in order to reach her destination. As each has been able to plot a complete route from origin to destination, the street names at each corner need only to be verified against the map to ensure that the correct street is selected. Here, the objective and path required to reach the objective is clearly defined. It is quite possible that the two shoppers will need to pass through the same corner, but subsequently choose different streets, in order to reach their destination. The decision as to which street to choose is comparatively easy, once the route has been planned. However, we can make it rather more complex by introducing a couple of changes. We first of all replace the maps with compasses, and secondly, we describe the location of each of their destinations in terms of landmarks. The shoppers are blindfolded and taken to arbitrary locations in the city. The first shopper is told that his destination is a clothes shop next to the Cathedral, which is due West from his current position. The second shopper is told that her destination is a newsagent near Trinity Bridge, which is due South. This time the

shoppers have only a reasonable idea of what they are looking for, but only directional signs about how to reach their respective destinations. At each corner, they have to decide which street is likely to bring them closer to their destination. As the city is not built on a grid, a street that initially looks promising, because it appears to be heading in the right direction, may turn out to be a dead-end or may ultimately tail off into the wrong direction. The shoppers may have to backtrack, revisiting the same corner several times in order to try out a different street. The paths of both shoppers may criss-cross several times, sometimes even trying the same street, hoping that there will be another corner which will bring them ever closer to their destinations.

Navigating through an information space should be an experience similar to that of the shoppers' when armed with a map and able to precisely work out a path from some arbitrary point of origin to a destination. Unfortunately, most browsing experiences in information spaces tend to be similar to the latter experience, partly due to the inappropriate organisation of information, but also because users do not always know what they require, and so would not be able to plan a route anyway. We vaguely know where we want to be, but we have only limited assistance to help us get there. Promising links often disappoint, whereas innocuous links occasionally deliver the goods.

How can the decision making process of choosing a link to follow be facilitated? We can, once again, provide an analogy to another, more pleasant, experience of shoppers orienting themselves in a city. The shoppers are, once again, given information about their respective destinations. As in the second example, they are given landmarks next to which they will find the shops. They are taken to different points of origin in the city, and are asked to find their destinations. This time, however, as they are trying find paths to their respective destinations, they may come across signposts which guide the way to landmarks. As they approach their destinations, these signposts will be available to guide them through the final stages of their journey. Signposts are typically present on corners, and so, with a bit of imagination, they can be considered similar to recommended links in a hypertextual information space. The shoppers still need to be able to arrive at a close enough location to their destination to find the signposts, and signposts may not be present at every corner, but users may reason or use their intuition during the decision making process. For instance, the Cathedral seeker may reason that Cathedrals are usually close to the city centre, and the bridge seeker may intuitively know that bridges are usually located over rivers, so this information can be used to get relatively close to their destinations. Once close enough, signposts will guide them the rest of the way, relieving the shoppers of having to decide which street to take at every corner when they are already reasonably close to the desired destination.

In a city, signposts are passive - they are either there or they are not. Signposts will give directions to all places of interest in the locality, regardless of whether the pedestrian is actually interested in that place of interest. Certainly, the city does not know in what each pedestrian is interested, and so the appearance of a signpost is not a reaction to a pedestrian's stated or implied interests. However, they are definitely useful aids for those who are interested in visiting places that are signposted. Is it possible to use a similar concept in a digital information space so that users are, at least some of the time, directed to the information they seek?

2.6 Assisting navigation in a digital information space

It is possible to assist the user in navigating through a digital information space by recommending links which are likely to lead to relevant material and by suppressing links which are unlikely to lead to relevant material. In both cases, the result will be to reduce the number of links which a user needs to consider. Of course, the user must be confident in the links chosen by the system, otherwise there is no apparent benefit - the user will still want to try the links that are suppressed or not recommended to ensure that nothing of importance is missed. Similarly, in a town, if signposts are pointing in the wrong direction, then pedestrians will quickly learn to distrust them. Another solution could be to assist users based on what prior users have done in the past. In social navigation systems [23], as long as the user has the same overall requirements as the majority, then the system will be successful. Otherwise, the success of the system depends on the granularity of the popular decisions and the overall ability of the system to correctly infer user interests. There needs to be some mechanism for identifying just what it is the user is looking for, suitably representing the domain so that its components can be compared with the user's interests, and subsequently tracing a path to the required information by recommending links for the user to follow.

The partial picture we have at the moment is of an unstructured information space. Meta-information enables the information space to be partitioned as a user browses, and also allows the same document to be used in different partitions. Links between documents are created to extend a partition as a user browses, reflecting relationships that exist between the document just visited and other documents that exist in the information space, given the nature or properties of the partition. If different users were to compare how the same document was presented to them, they would see that its apparent location in the information space is different, as links presented to one user are not necessarily presented to the other. Whenever a user can be automatically guided to the required information, links are recommended.

Although the basic content of a document is constant, when it is presented to a user it is influenced by being encapsulated by the extended partition. This influence results in a change to the document. The scope of the change that has been suggested so far is that how the document is linked depends on which other documents in the information space are relevant in a particular partition, as determined by the document's meta-information. The scope can be further reaching to include modifications to the actual information presented in the document. For example, data can be presented as a pie-chart or a graph, depending on which is more applicable in the partition. Paragraphs of text may be hidden from the user, or glossary terms may be expanded. The potential for changes is limitless, but is always determined by the manner in which the information space is partitioned.

2.7 Summary

This chapter has identified some of the problems facing a user seeking information in a digital information space. We have used a simple analogy of a person recently arrived in a new town looking for basic services, such as a newsagent, to show that part of the navigational problem inherent in information spaces is a lack of commonly accepted and known methods (such as structure and terms of reference) to orient users within the environment. In large to massive information spaces, the problem is compounded because of the increased choices and lack of geographical limits, reducing our ability to exhaustively browse all the options whereas simultaneously lulling us into believing that precisely the information we seek definitely exists. A digital information space is also usually featureless - we cannot tell how close we are to the information we seek, because there is often no indication about where things are in relation to each other.

We have argued in favour of meta-information to provide features within a digital information space. Meta-information allows information to be related according to some definition, as well as allowing the same information to be re-used. In addition, by partitioning an information space the user is presented with a smaller information space within which to make decisions. The decision-making problems facing a user are exemplified using the analogy of a shopper locating a shop in town. Just as a shopper may use landmarks, and signposts to those landmarks, to locate their ultimate destination, a user browsing through a digital information space may be assisted by the presence of recommended links which are considered likely to lead the user to relevant information. Unlike signposts in a town, recommended links are displayed to users only if they are relevant.

HyperContext is a new framework for hypertext in which users are supported while browsing or searching for information. As a user browses through a HyperContext

hypertext, a partition is created. Whenever a link is followed, the destination document is interpreted in context prior to being included in the context session. Meta-information describes the interpretation of a document and determines which links will be displayed to the user and what destinations each link will have.