Dieberger, Andreas, and Andrew U. Frank. "A City Metaphor to Support Navigation in Complex Information Spaces." *Journal of Visual Languages and Computing* 9 (1998): 597-622.

A city metaphor for supporting navigation in complex information spaces

Andreas Dieberger, Emory University, email: andreas.dieberger@acm.org
Andrew U. Frank, Technical University of Vienna, email: frank@geoinfo.tuwien.ac.at
published in *Journal of Visual Languages and Computing, 1998 (9), pp. 597-622*
Note: Sorry about the bad quality of the figures -- we temporarily misplaced the original files

Abstract

A major problem in modern information systems is to locate information and to re-find information one has seen before. Systems like the Word-Wide Web are heavily interlinked but do not show structures that help users to navigate the information it contains. The use of appropriate navigation metaphors can help to make the structure of modern information systems easier to understand and therefore easier to use.

We propose a conceptual user interface metaphor based on the structure of a city. Cities are very complex spatial environments and people know how to get information, how to reach certain locations in a city, and how to make use of the available infrastructure etc. Cities provide a rich set of navigational infrastructure that lends itself to creating sub-metaphors for navigational tools. A city metaphor makes this existing knowledge about a structured environment available to the user of a computerized information system.

We first focus on several properties of future user interfaces (or user interface metaphors) that will distinguish them from current systems, like the richness of information or the use of visualizations to show the structure of information spaces. We also describe the strengths and problems of spatial user interface metaphors. Then we describe the structure of the information city metaphor, its structuring and navigation metaphors and what we see as its main advantages and problems. We further describe a few scenarios of how an Information City might work. Finally we compare implementing this metaphor using either a textual or graphical virtual environment or a combination.

1. Introduction

In modern information systems the main problem is not any more to get more data into the system, but instead how to find and later how to re-find information one has seen before. Heavily interlinked systems, for example many World-Wide Web sites, are difficult to navigate because of their lack of apparent structure. Appropriate *navigation metaphors* can make the structure of information systems easier to understand and therefore easier to use.

User interfaces commonly are based on metaphors that help understanding the system in terms of objects the users know already. Many of these metaphors are based on real world objects, like folders, notebooks, trash cans and so forth. One class of real life metaphors is spatial metaphors. They exploit the extraordinary human ability to organize objects in space, to recall and reason about their locations and many other space related cognitive abilities [32]. Probably the best known example of a spatial metaphor is the desktop metaphor, now widely used on personal computer systems. For an early account of the desktop metaphor see [52]. It is well suited for managing files, yet the basic desktop metaphor (originally designed for several hundred documents) reached its limits with file systems that contain tens of thousands of files, which is a typical number for most personal computer systems today. For even bigger systems, like the World-Wide Web, a desktop metaphor is totally inadequate.

The desktop space can be extended to an office or room metaphor and further to building and city metaphors. Each of these metaphors has its own merits and limitations. An important advantage of building and city metaphors is that they define several levels of enclosed spaces. Therefore, they allow to create structures with many levels of security. They also support interactions between multiple concurrent users in the environment.

In this paper we describe a user interface metaphor based on the structure of a city. This metaphor, called the *Information City* is a conceptual system that focuses on the issue of navigation and the recall of places, but also provides interesting features for issues of security, privacy, visualization and so forth.

The Information City is not a metaphor meant to be implemented straight from our specifications. Instead it defines an ontology of spaces and connections that we think is useful for talking about systems of spatial metaphors and how they interrelate.

Cities are very complex spatial environments and people are used to navigate cities. They know how to get information and how to reach certain locations, how to use the infrastructure and so forth. A city metaphor makes this existing knowledge about a structured environment available to the user of a computerized information system. The navigational infrastructure available in most cities allows creating sub-metaphors for navigational tools.

Cities traditionally are spaces where people navigate collaboratively. They guide each other, they point out landmarks, they give (sometimes incomplete) route descriptions and they interact directly with each other and with objects in space. These properties of the dynamic city environment and its navigational infrastructure are the qualities we try to support in our Information City metaphor. The Information City we envision is a tool to either communicate structure of an information space to the user or to explicitly create structure in an unstructured information domain.

Structure of this paper

In section 2 we describe the problem of navigation in complex information spaces. We focus on the challenge of communicating the structure of an information space to a user. Navigation is defined as a mapping from such a structure to navigation activities. We argue that navigation can be easy only if the structure of the space is understandable.

In section 3 we describe spatial user interface metaphors and how they can help in navigation by making structures explicit. We describe advantages and disadvantages of spatialization and that future spatial metaphors will differ from earlier systems both in their use of history information and rich visualizations as well as in their use of magic features.

Section 4 outlines architectural and city-planning concepts that we used in the design of the Information City. City-planners have a good understanding of how to structure large architectural spaces to make them easy to use and to navigate. Our taxonomy of city elements is strongly influenced by Kevin Lynch’s model of people’s conceptual model of the city environment. We further review related metaphors in literature.

In section 5 we define the Information City by describing its elements, their main properties and how they interact. The Information City is a system of metaphors that allows us to create complex information spaces using a structural framework based on the city. The city elements are split into structural elements, navigation tools, information providers, and magic features. In section 6 we describe a scenario using an Information City.

In section 7 we discuss implementation issues for an Information City using either a textual or a graphical virtual environment. In particular we discuss the design of magic features and the visualization of usage information. We further outline possible strategies for how to start building a city.

Section 8 draws conclusions from the previous sections, points out unresolved issues and summarizes the paper.

2. Navigation in Complex Information Spaces

Appropriate user interface metaphors facilitate understanding the structure of an information space, which helps users to find the information they look for. This structure stems from the information domain and is communicated via the user interface (metaphor). The user interface metaphor itself imposes additional structure on the information space.

The term *information space* hints at a spatial conceptualization of the metaphor. We believe that spatial user interface metaphors have advantages for navigational and organizational tasks. There is a strong relationship between spatial metaphors and information visualization: the visualization communicates the structure of the information space so the user can easily navigate it.

2.1. Movement in Information Spaces

Navigation is the process of mapping the perceived information structure to activities for accessing information. Navigation is possible only when structure is communicated to the user. A visualization of the information space in the most general sense thus is a prerequisite for navigation in that space. Navigation tries to answer the following questions:

- "Is there 'a piece of information' with the property X?"
- "How can I access this information?" or "How do I get to this information?"
- "Where am I now in relation to ...?"

The spatial formulation of these questions allows us to understand every information environment as an information *space*. If the user interface supports navigation by clearly showing the information space's structure this spatial concept becomes even more explicit.

Note that users do not automatically inhabit a position in space. Also there is not always a history of where the user has been. The availability of such concepts depends on the user interface metaphor and therefore on the visualization of the information space.

Hypertext

Most information in modern computer systems is interlinked and exhibits hypertext character [45]. Hypertext is a space constructed of information nodes and links connecting those nodes. Users navigate the hypertext network using those directed links.

The basic principle of hypertext is *non-linearity*. Where most types of information have a typical order in which they should be apprehended this is not necessarily true for hypertexts. Hypertext reading has been called *associative reading*, as the user determines the reading order dynamically in an associative way.

Users easily loose orientation even in small hypertexts. This orientation problem is called the *getting lost in hyperspace* problem. As in every other information structure, hypertext navigation relies on communicating the structure of the information space to the user. Most hypertexts lack a perceivable structure, which is the main reason for the navigation problems in these systems. However navigation problems are not inherent to the hypertext concept but are caused by a user interface that fails to communicate the structure of the information to the user.
To support navigation in large information spaces, including hypertext, we need to provide tools that make the information structure apparent and that help answering the core navigation questions reliably and efficiently [14].

2.2. Information Visualization

One possibility to communicate the structure of an information space is to visualize it using a graphical representation. Structure visualizations of this kind are one of the main working areas of the field of information visualization.

Examples of structure visualizations are the Information Visualizer [8], the Navigational View Builder [43], the Narcissus system [23], or the visualizations for the Hyper-G system [2]. Visualizations of information spaces use spatial metaphors and arrange their elements in 2D or 3D space. The visualizations express meaning through proximity, alignment, size and color cues etc. Without this perceived meaning a spatial arrangement is as useful as a graph with an unlabeled chart axis.

It is important to point out that although the underlying metaphor is not always made explicit in a visualization, it does influences the design of the visual representation.

It is generally assumed that spatial visualizations of information structures help users to learn the structure of an information space. However Poblete cautions us in [47] that simpler and flatter representation of information structures sometimes are learnt easier than more advanced 3D visualizations. This shows that the silently accepted credo of "3D is better than 2D should be reassessed. We expect to see more such critical evaluations of information visualizations in the near future.

An important aspect of visualizations aiding in navigation is that they shall not overwhelm the user with information. In particular for 'view navigation', Furnas showed that it is ideal to show only small views (a relatively small number of choices), that the number of navigation steps is not too large and that the route to any target must be discoverable [21].

3. Spatial User Interface Metaphors

In this section we focus on advantages and disadvantages of spatial interface metaphors. Their main problem is that they do not scale well. This problem occurs both for the visualization of the metaphor as well as for the navigation in the information space.

Other issues of interest are the amount of realism used in the metaphor's representation, how to represent a history of use, and *magic features*. A magic feature improves the efficiency of a metaphor by breaking it.

The usefulness of spatial metaphors is intuitively understandable as well as proven in several studies. Two studies that reported independently that users like to organize information spatially are summarized in [5]. Two other important, but earlier papers on this topic are [31, 39]. However spatial metaphors do have their problems which need to be addressed to make them the basis of an information system.

3.1. Advantages of spatialization

An advantage of spatial organization schemes is their initial familiarity for users. Most people organize objects spatially -- be it in folders on a shelf, or in piles on a desk. Most people find it also easy to remember approximate locations for objects they placed in space.

Humans have strongly developed spatial cognitive abilities and human language provides many constructs to describe space and locations in space. This enables people to use even incomplete or partly incorrect descriptions of locations and routes.

A spatial arrangements can express relationships between objects metaphorically that are difficult to describe formally. Spatial language can describe complex relationships and even leave room for ambiguity. Similarly most people easily understand the arrangement of objects in a visual representation as showing implicit relationships.

A good example for how complex spatial relationships can be is the perceived structure of a city as overlapping neighborhoods or districts. City areas do not form a strict hierarchy but rather a structure of partial containment [1].

Similarly information objects can be arranged in space as a visual expression of complex document relationships. Spatial arrangements based on a spatial metaphor thus are a useful tool to visually express complex relationships between objects.

Spatial metaphors further provide source domains for metaphors representing access paths, enclosure, forbidden access etc. which can be the basis for navigation tools in information spaces.

3.2. Disadvantages of spatialization -- Scaling up

The major problem of spatial metaphors is that they do no scale up well. Scaling is problematic both for the visualization of the metaphor (how to represent very large spaces) as well as for the navigation (how to access a location that is remote in that space). The navigation aspect of the scale problem will be discussed in section 3.4.

The desktop metaphor is a good example for the scaling problems. The file system of a modern desktop computer typically is so huge that it is impossible to represent many of it's directories using the desktop metaphor. Therefore users tend to choose a list view rather than a symbolic (spatial) representation to view their directories.

In this example the close relationship of spatial metaphors and information visualization becomes apparent. A more elaborate visualization of the file space can solve this problem: One of the features of the Information Visualizer system is to support the representation of a very large file system as a cone tree, a 3D visualization of an entire file system [8].

The main challenge for visualizing structures is probably to provide a *global view*. Examples for system that are steps in this direction are the table lens system [50], the Perspective Wall of the Information Visualizer [38], or the Information Mural [29]. These systems proved to be fairly successful even for quite large numbers of objects. Another example is to simply add a third dimension to the desktop metaphor and to allow objects to be pushed back in space. This has the disadvantage that close objects may obstruct objects farther away so that users don't get a complete view of all objects.

Many visualizations eliminate scroll bars. Instead of using a window onto an information space (which shows only a small part of the space) they provide a global overview and therefore maintain the context in which information is located. The user can zoom in on a small detail but the global view of the space is available all the time. These systems still have to prove they scale up significantly better than the plain desktop for general purpose applications.

3.3. The necessary amount of realism

An important issue in the visualization of a spatial metaphor is the amount of realism in the representation. Spatial filing is a useful way to organize a large number of diverse objects exactly because these objects are diverse. Spatial filing does not work well when the objects look too similar. Representing a space and its contents realistically helps creating a space with diverse objects. But how much realism is really necessary? Can a metaphor be too realistic?
While a detailed recreation of a metaphor’s source domain helps transferring preexisting knowledge to the target domain, we often need to deviate from the metaphors source domain to eliminate usage modes that do not fit the target domain.

When assessing the usability of the overall metaphor we have to find an adequate balance of supporting various aspects of usability. In particular we often have to find the right balance between learnability and efficiency [44]. In many cases this balance determines the level of realism necessary for the visualization. A more realistic representation makes the metaphor easier to learn, but might prevent the inclusion of efficiency enhancing features.

Consider the example of an electronic book and the task of turning a page [30]. This could be realized by actually grabbing a page with the mouse pointer and pulling it to the other side, which is much closer to the source domain than a "next page" button. Users probably would find this feature very tedious because forward and backward buttons are much more effective and useful on a computer system.

As additional example consider a file system. While text-only interfaces use only a file name to distinguish between objects, graphical user interfaces use a combination of icons and filenames. Although a file then has the properties location, name and file type (maybe even color) it is still necessary to scan folders with many documents of the same type by reading all file names. This is especially true when the spatial arrangement is often reorganized. Users then revert to alphabetical or temporal lists and the location information of the spatial metaphor is useless or even slows users down.

To further enrich file systems it is possible to use icons showing the content of the document (proxies) [25], color cues for the age or size of the document[53]. Other authors proposed using different communication channels. For example the Sonic Finder [42] used sound cues to convey information about a document's size and with the availability of haptic output devices it is very possible that documents soon might also 'feel' different according to their contents. All of these approaches have one thing in common: they try to enrich the information space described by the user interface.

While these enrichments are very helpful for the user others are more of a nuisance. For example early spatial metaphors tried to extend the desktop to a room metaphor by showing typical office objects arranged in the representation of a room. The main object in these metaphors was a desk located in the middle of the room. Unfortunately such a view will represent the desk surface at an oblique angle and make it impossible to discern documents on it. Essentially these systems provide the user with a fixed location in space. The overly realistic recreation of the room metaphor lead to a representational problem that the software could not compensate, as the graphics capabilities of computers then were not sufficient to allow users to freely move in space.

Users must not be tied to a fixed location in a spatial metaphor. At least the represented objects must be movable. An example for a solution is to represent the documents on the desk at the oblique angle and to allow the user to pick them up to seem them in plain view. This approach was realized in the CD-based games MYST and Riven and also in the Web book system [9]. Still this does not really solve the problem, as the user will have to pick up one document after the other to see their contents. One of the advantages of spatial metaphors is that users can perceive an information space as a whole, taking as much information in as possible in one view.

A very realistic representation can have advantages when it is used as information carrier. A good example for this concept is the idea to visualize usage information as *read wear*, first described in [24]. The basic idea of read wear is that digital objects *wear out* like physical objects. Visualizing wear thus hints at objects that are used especially often and therefore may be of special interest to the user population. Recent work on group memories, social navigation [18, 60] and recommender systems [49] sometimes combines the ideas of read wear with a voting scheme to consider also perceived importance. For examples see [24, 40, 48, 51].

3.4. Navigation and Magic features

Spatial organization is no universal cure, even with a representation that scales up well. Users need a means to move in the information space. As we described in section 2 navigation is the mapping from a representation of the information space to actions to move in this space. A good example for such a mapping is the concept of sites, modes and trails as described in [46]. An example for a computer model of the navigation process focussing on the mapping to navigational activities is [33]. While these references consider conventional space, it can be desirable to provide navigation features that go beyond conventional space.

In conventional space movement causes effort proportional to the distance and navigational means available. However to effectively use a spatial organization scheme shortcuts through space are essential. These shortcuts may break this relationship between distance and effort to travel. They appear as something that lies outside the underlying metaphor and therefore we call them *magic features* [14, 15].

As mentioned in the previous section the creation of a usable interface metaphor requires finding the adequate balance between learnability and efficiency. While a focus on realism in the interface helps users to learn the metaphor, magic features provide the necessary shortcuts to make the metaphor efficient. They are necessary to increase the overall usability of the system.

Magic features need to be designed in a way that they do not compromise the learnability of the system. Therefore magic features have to be used in a controlled and restricted manner -- otherwise the spatial metaphor falls apart and becomes confusing for the user.

Controlled and very limited breaking of the metaphor seems to be a principle in all successful spatial systems. An example is the Macintosh desktop metaphor where a folder can contain another folder (which would not work with real file folders). Folders also can be represented by aliases, and an alias of a folder can be located inside that same folder. Using aliases it is possible to create a one-step shortcut to a document or folder that otherwise would be buried deeply in the folder hierarchy. Similarly, a search function in the desktop metaphors is also a magic feature: it transports the user ignoring regular access paths.

Possible magic features in a city metaphor are a doorway that transports a user a far distance like an alias (teleportation), or a search function.

Because magic features are outside the overall system metaphor users cannot draw on the metaphor's source domain to understand the feature. Even worse, they might assume the magic feature is indeed part of the metaphor and are then unable to understand the magic feature’s working.

Magic features have to be designed to look different enough from the rest of the system so that users recognize them as something special. Spatial metaphors with magic features can reduce the navigation problems in very large systems. While a strict spatial metaphor does not scale well for navigation tasks, magic features provide shortcuts through the system that can structure a large information space into smaller consistent regions and thus make navigation effective even in very large systems. A careful design of magic features allows users to use the spatial framework provided by the spatial metaphor and still understand the shortcuts provided [16, 32].

4. The Information City - motivation and basic concepts

The *Information City* is a conceptual spatial user interface metaphor for large information spaces. It is based on structures found in real cities, on knowledge of city-planning and on how people learn such environments. The city is a rich environment with which humans have extensive experience. It is an excellent source for a metaphor because it is extensible and can be navigated using commonly available infrastructure. In this section we explain our choice of a city metaphor, describe the Information City’s fundamental structures and discuss other systems based on city metaphors.

4.1 Why use a city structure?

From our discussion of spatial metaphors it may seem reasonable to extend the relatively limited desktop space to a larger space, for example a city. The main argument for the city metaphor is that starting with the complexity of a city the metaphor provides dynamic structures that scale up comparatively well and also provides useful navigation tools to cope with the space's complexity. Another argument for the city is that people tend to know only small parts of cities well and easily learn how to navigate between these parts. They also know strategies to navigate effectively even in unknown parts of the city.

The structure of cities is an additional benefit. Although cities are hierarchical in a way, people's mental representations of cities contain overlapping elements. These hierarchical and overlapping structures are useful for sub-metaphors that describe not only a hierarchical structure but also more complex relationships of containment.

Many of the city elements can serve as container metaphors [34]. Examples are the district, the neighborhood, the block, the building and so forth. Several of these elements have strong boundaries (for example the building or the room). They are ideal sources for metaphors that describe strong encapsulation and access control.

People are used to navigate city structures. They know which navigation infrastructure to expect and which transportation means to use for which navigational task. People also have learnt social protocols for various city elements, which control the interaction with other users in these areas. The city therefore is a social space as well. These topics will not be covered in detail in this paper. For more information on this aspect refer to [14].

A final advantage of a city structure is that information rich city environments are relatively easy to learn if frequently used. Tauscher and Greenberg report in [54] that 58% of individual's Web pages accesses are visits to previously seen pages. Learnability of the information structure is one of our main goals. Therefore Tauscher and Greenberg's study encourages us to choose a metaphor for representing information systems that is especially useful for repeated visits -- especially for a system like the Web.

4.2. Elements of the city environment

Cities are grown organic wholes. Their structures develop out of the needs of their users and from the interaction of many people. People acquire a mental image of the city environment based on these grown structures. In this section we describe the main elements comprising such a mental representation. These elements are not necessarily identifiable with objects in the environment. They rather have to be understood as elements in the mental representation people construct of the city environment.

Kevin Lynch's study "The image of the city" [37] describes five major elements in the city image. They are the *Node*, the *Path*, the *Edge*, the *District* and the *Landmark,* but these are not the only city elements conceivable. City elements are not always clear-cut. Their fuzziness stems from the fact that the city environment is not strictly hierarchical.

A *node* is a point-like element in the perceived image of the city environment. In the structure of the city a node can be the crossing of two major streets or other linear elements. For navigation purposes people often refer to such nodes when describing their location.

A *path* is one of the two linear elements in the city environment. A path is a mental concept describing how to get from a location A to a location B. An example is the path from home to the office. Paths can be major streets in the environment but they do not have to as they are a more general concept.

Edges represent borders or visual separations in the environment. While most edges are sharp also smooth transitions can be perceived as edges. Also this linear element may co-occur with physical linear elements in the city, but it doesn't need to. Consider a highway leading through a city. It is a strong separating element between two areas (edge) but can be perceived also as a path -- depending on the point of view and the actual context.

Districts are areas containing objects with a common element or character. The common character may be in the style of buildings, the prominent use of buildings, or another not necessarily visible aspect. Districts should not be confused with bureaucratic districts, as they are no clear-cut entities. While a bureaucratic district has a sharp boundary (often exactly in the middle of a street), the district as a city element can be a fuzzy concept. As alternative we could use the word *neighborhood*. Districts show clean boundaries when they are bordered by an edge. This is the case in cities with historic centers where city walls provide a highly visible edge.

Landmarks are elements showing *unmistakable form*. Landmarks must be discernible from all other instances of similar objects in the environment. They are -- like nodes -- conceptually small elements. Landmarks are of special importance when giving directions since they are easily recognizable.

A well-designed real city provides a well-balanced mix of the five city elements, see also [36]. This allows users to easily learn paths, to describe and remember routes and locations. It also may give locations in the city a feeling of *place* which provides context for objects nearby and a framework for social interaction [20].

4.3. Additional richness of the city environment

The city elements described in the previous section provide the essential structural elements for a city metaphor to be usable. The city metaphor can adopt additional city elements and characteristics to enrich the bare city with additional information, much like we can enrich a graphical user interface as described in section 3.3.

A city that consists only of similar blocks with little or no differentiation is difficult to use even when plenty of structural elements (like paths and landmarks) are available. Finding a particular building in a district would still involve looking at each of them in sequence -- just like scanning a list of file names. If buildings look differently, possibly giving an indication of their contents (like proxies), age (building style), and use (using read wear) then finding a certain building in this environment will be much easier.

As the user's needs change also the visualization of additional information may change according to the task at hand. This change provides several different *views* of the same environment, each of which can be optimized for certain navigational tasks [19].

For the city metaphor to scale up we need to depart from a strict Euclidean space in some situations and introduce magic features. These magic features are also an enrichment of the environment because they go beyond the city structure and beyond the source domain of the city proper.

In the context of the city magic features can have the form of magic portals that transport a user to distant locations. Besides such navigation features there is the possibility of magic in structures. A possibility is a building or a room that contains an area that is larger that itself. Such a feature, the envelope, will be described in the next chapter.

Note that magic features are an addition to the environment that should be used sparingly. Filling the city environment with too many magic portals, for example, destroys the specialty value of these elements. Users will not see them as an addition to the structural elements of the city but as the main structure. The result would be a hypertext of buildings without the structural benefits of the city metaphor. However, we do believe that magic is an important ingredient to improve the scaling behavior of the city structure and make it more effective.

4.4 Other city-like metaphors

Studying other city-like metaphors in the literature and on the Word-Wide Web made us realize that most of these systems do not incorporate the vast amount of knowledge already existing in the fields of architecture and city planning. Also most of these systems have no consistent building concept -- problems we try to avoid in our information city metaphor.

Several spatial metaphors have been based on the structures found in buildings and cities. Examples are Magic Cap or the Xerox Rooms systems [22].

Further extensions can be seen in systems like the Digital City (http://www.viper.net/fun/dc), the WebWorld (a discontinued system on the World-Wide Web), the Paxton virtual city (http://www.eolas.co.uk/mellanta/fd/) or Apples e-World, all of which make use of a village or city metaphor. Recent systems extend the spatial metaphor to a multi-user VR world to meet people and to chat, like the AlphaChat system (http://www.worlds.net/) to name just one example.

Most of these systems present the user with a relatively small space in which she can arrange objects according to her wishes. They tend to scale badly as their city metaphors merely add metaphorical sugar to an existing system. In systems designed mainly as meeting spaces (for example WorldsChat) navigation of large amounts of information is a minor concern.

Many existing city metaphors introduce mainly a new term for the folder concept. They seldom explore novel navigation tools and they almost never leave the concept of strict hierarchies. Furthermore the interactive aspects of cities is lost and if the systems use a graphical representation it is often more a nuisance than a useful tool. In most of these environments buildings and rooms are simply containers and arranged in whatever fashion.

An example for a system of this kind is the Paxton virtual city. It uses buildings to group links to information on the World-Wide Web but there is no concept behind the decision where certain links are located so that every building can serve as a container for whatever information. Paxton does a good job to demonstrate what can be done with the WWW but it does not provide adequate structuring and navigation tools for large information spaces.

The WebWorld system (not on-line any more) was similar to Paxton. It did not provide a container structures but rather a large, user modifiable landscape into which people could place their objects (buildings). Although structurally quite simple the WebWorld provided magic features for navigation: users could create their own WebWorlds and create links to them through *portals*.

Recently there is growing interest in city like metaphors through advances in information visualization. An example is the Information Space described in the Starfire concept [57]. Another promising idea is Rob Ingram's work that uses Kevin Lynch's city elements as the basis of an information visualization system [26-28].

Also of interest are the visualizations created for the HyperWave system, a second generation distributed hypertext system that shows many advantages over the architecture of the World-Wide Web [2, 3]. For information on Hyper-G (now HyperWave) see [41].

Magic features seem to occur more commonly in systems designed to support virtual meetings, for example in the Worlds Chat system. Also many textual virtual environments make use of city-like metaphors. These systems do not focus on navigation of a consistent information structure which maybe makes it easier to get away with magic features.

5. Description of the Information City Metaphor

In this section we describe the *Information City* metaphor. Because of space constraints the ontology described here is only a part of the metaphor. For more details see [14]. We start with the structural framework as the basis for the navigation tools. Information content is provided though information providing elements. To improve the scaling properties of the city we introduce several magic features. In an additional section we discuss issues of actually creating an Information City.

5.1. Topology and Overall Structure

The topology of the Information City is based on generalizations of Lynch's three major elements: container, landmark and path. The Information City consists basically of a collection of containing elements that are associated with at least one landmark.

The largest containing element is the city itself, but there is no restriction to only one city. The next smaller element is the *district*. Inside districts there may be sub-districts that consist of buildings. Buildings contain floors and rooms -- according to the structure of the information organized within the container.

Each container can be considered a complete subspace that does not have to adhere to the general city framework but may be organized using a different metaphor, should that be adequate.

5.2. Lesser Structural Elements

When closing in on the city, major navigational paths that separate large districts come into view first. As districts contain related objects, the separation into districts can be highlighted in the visualization. Getting closer brings lower order paths and smaller containers into view. Buildings and small-scale areas can be recognized. Finally, single buildings, small-area landmarks, architectural properties of single buildings, and specialized buildings, like subway stops, can be discerned.

5.2.1. Containers

- A building is a container for information or infrastructure in the Information City. Buildings have a unique address and show their accessibility using doors.
- Landmarks are special non-access or public access buildings. In the first case, the building has only landmark function, e.g. a clear vista at the end of a path. This case is useful if landmarks are placed at the center of a cluster of related documents if there is no object available in that location to serve as landmark. In a graphical realization of the Information City landmarks can be seen from far away. Major landmarks should be higher than most buildings to provide orienteering aids for users flying over the city.
- *Rooms* are containers inside buildings. Their walls may contain doors or windows to access other rooms or the outside. Rooms show their accessibility through doors. Like other containers, they can contain a non-spatial metaphor, like a view-screen.

5.2.2. Navigation Infrastructure

- *Paths* connect two locations in the city. Therefore, they have starting and ending points that should coincide with landmarks. As in the real city, a path is a continuous element of the Information City. Paths outside buildings are visualized as streets or roads.
- Intersections of paths are *squares*. Large squares are major elements in the city. They contain stops of public transport systems and they can contain landmarks. Squares correspond to the nodes in Lynch's city elements.
- Lines are linear elements that don't adhere to the Euclidean space concept. Transportation, for instance the subway, travels on lines. Lines are connected to the city environment in a few distinct locations. Paths and lines should be visualized differently.

Inside buildings (containers) similar elements occur as described above, however, they have different names and sometimes, slightly different functionality. They are the *hallway*, the *lobby* and the *elevator*. The elevator behaves like a line and also provides access to navigational infrastructure outside the building.

5.2.3. Connections and Separations

Connections and separations are conceptually similar structures, however separations provide access control. Normal navigation structures do not restrict movement.

- Doors connect locations inside a building to indoor locations or to the outside. This state of doors (open, closed, locked, ...) shows the accessibility of the corresponding room. Open doors can be looked into. This "preview" shows at least if the door leads to a room or to an envelope.
- *Windows* are similar to doors, but are mostly located in facades. Although they can be used as departure points and destinations for flying, they are not meant as major paths but to look through. A special type of window is the magic window. Windows are a sort of shortcut to rooms inside a building.

5.3. Navigation Tools

We distinguish between transportation and navigation tools. Although both of these are navigation in the traditional sense, we consider transportation to be a more passive form of navigation where the user is moved, whereas navigation is an active process. For example, consider the difference between taking a taxi and driving a car. In a taxi, the user instructs the driver as to the desired destination, and then sits back for the ride. On the other hand, driving her own car, the user has to choose the route and steer the vehicle herself. There are several types of navigation and transport metaphors according to the navigation task.

Navigation

- *Walking* is navigation for short distances. It uses paths, squares and all open access structures. Walking can be half-automatic when an address (a link) close by is selected or the user decides to follow a "red carpet".
- *Driving* is a metaphor for fast walking and for covering medium distances.
- *Flying* is used for long distance navigation.

Transportation

- *Taxis* are like cars, but are not constantly controlled by the user. They can be summoned from anywhere and are able to navigate using incomplete information. They can even provide a guided tour.
- The *subway* provides long-distance transport in the city. It has a set of predefined stops which always coincide with major landmarks. Leaving the subway at those landmarks either places the user in the main lobby or in front of the landmark. The travel time gives a rough indication of the distance traveled. Subways do not show the environment traveled. Subways can be left only at predefined stops but a temporary subway stop can be summoned anywhere outside buildings to enter the subway. Inside buildings elevators provide a connection to the subway. The use of the elevator heightens the user's awareness that she is leaving the building.

  The subway *tunnels* through the city space. It connects distant points without traveling through all locations in-between and thus, travels in a different space than a *walking* user. The subway's space concept is based on connectedness. This movement of subways in a *somewhat detached space* was observed already by Kevin Lynch in [37].

Depending on the distance traveled, users apply different navigation and transportation tools, which need to be realized differently. A classification according to distance decides on the visual representation of the navigational activity and its *enactment* [6, 35]. Correct enactment helps users to better understand from where to where they are navigating. The enactment must provide the necessary feedback so that the user does not get disoriented in the navigation process.

5.4. Information Providers

Information providers are the information-carrying entities of the city. Some of them are specialized structural elements.

Walls and Signs

- The *facade* of a building is not only an wall, but it provides information about the contents of a building. Facades should show read wear to indicate usage use and visual cues to indicate the content of the building.
- *Information walls* present information and also show read wear. They can link to other information walls using the tramway or red carpet metaphors or though magic windows embedded in them. These links should also show their usage. Information walls cannot be moved.
- *Signs* are small information walls which provide no linking.

Movable Information Objects

- Removable information objects are based on source domains like writing pads, books, or business cards, and contain information that is not strictly associated to a location. They may be available as single objects or they may be provided by information dispensers that create an unlimited number of information objects. The dispenser is associated with a fixed location in the city, but the information object itself is movable. Examples of metaphors for dispensers include a newspaper vending machine or an information kiosk.

5.5. Magic Features

The city metaphor is usable without magic features, but we believe that it scales well only when magic features are added. In this section, we describe a structuring and a navigation feature. Most navigation and transportation tools (see Section 5.3.) show magic properties but are grouped with the other navigation and transportation tools.

- *Envelopes* are special rooms. They are outside of the Euclidean space concept of the city. Like rooms, they are accessed through doors, but they give access to very different structures or even another Information City. A city contained in an envelope can be used to represent archived data. The contents of an envelope may actually be larger than the envelope seen from the outside. A transition to an envelope has to be clearly enacted as a magic feature.

  Leaving the envelope transfers the user back in front of the door leading to the envelope. Such a transition has to reestablish context for the user e.g. by placing the user high above the city and zooming into the location she occupied before entering the envelope.
- Magic windows (or teleporters) provide a direct connection to other parts of the building, another location in the city, or a view-only connection of another location. They are a special case of the information wall. Magic doors and windows have to clearly show that they are magic features and, as something "special", they should be used only sparingly.

6. Scenarios

In this section we describe a short scenario of how an Information City might work. As we mentioned already the Information City doesn’t have to be represented as a graphical construct that looks like a real city at all. By using a scenario we can avoid planting a specific realization of the city in the reader’s mind. In the scenario we stress features that distinguish the Information City from other city metaphors, in particular social navigation tools, read-wear, and magic features.

6.1. Asking for the way and using read wear

A user enters the city searching for information on the use of "spatial metaphors in information visualization". She might first fly over the city for a moment, studying it's layout and decides to research in the computer science district. Entering the computer science district she notices fast what a busy and fast-moving area this is. There are plenty of kiosks around offering access to latest news in various areas of computer science, each of the news messages with a link to the corresponding sub-district or district, and showing how often is has been accessed. (Figure 1)

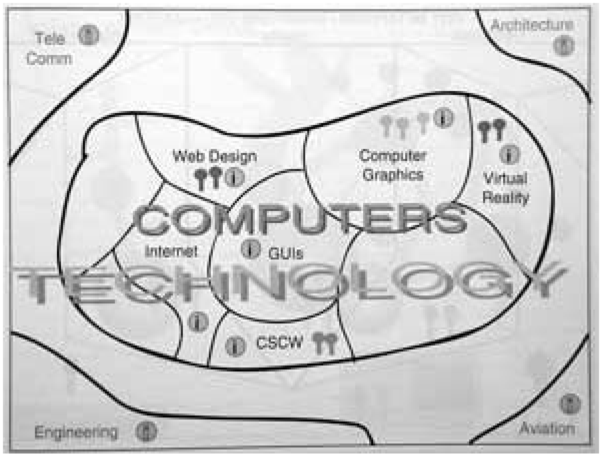

Figure1: Talking about information structures in city terms does not mean the structure has to look like a real city. This sketch shows how a district could be rendered as islands with sub-districts containing related information. Each sub-district does contain a landmark and indicates information like the number of users present, read wear etc.

She notices the kiosk for virtual reality news which might be a good start, but she didn’t want to limit her search to 3D visualizations. As she doesn't see a kiosk that better fits her search goal she approaches a group of avatars and asks for guidance. Several of them points out that all spatial metaphors and visualizations are near the virtual reality area and point into a certain direction. With such a strong recommendation the user decides to look at the virtual reality area after all. She walks in the indicated direction and the topics she encounters change gradually from general computer science to graphics, visualization, information visualization, and then virtual reality.

On the VR square there indeed is a building dedicated to spatial metaphors. The facade shows that this is a less frequently used building -- especially in comparison to the VR games building. This explains also why there was no note on the central computer science square. Entering the building she notices a highlighted hallway leading to currently hot topics. Walking along this hallway she eventually reaches a room about city metaphors that indicates it is occupied by users. She enters it and finds a group of avatars discussing city metaphors. (Figure 2)

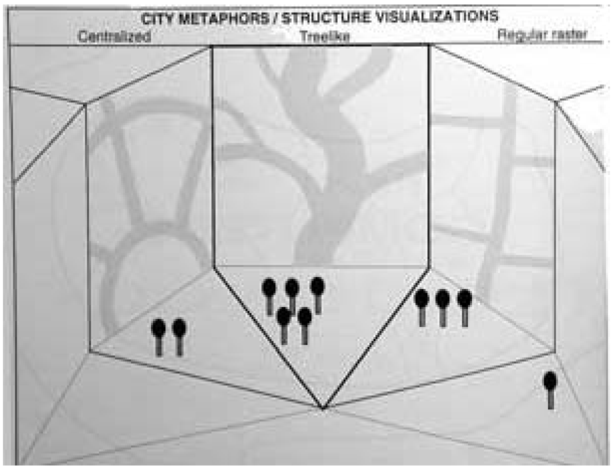

Figure2: A possible representation of a room serving group discussions. Several documents discussed are presented on different walls. The discussion can focus on one particular wall. The visualization can represent 'where people stand' in the discussion by representing their avatars in the appropriate sector of the room.

6.2. Envelopes, and magic features

Our user soon realizes that the discussion in the room is about systems she has never heard of and asks the group about a repository of information of spatial metaphors. One of the avatars points out that she is currently working on such a repository. She invites our user to have a look and hands her a business card with a magic window on it.

The magic window shows a locked room in the same building. It contains several view screens and envelopes. She activates the magic window to be transported to the spatial metaphors repository room. The view screens run demos of systems using spatial metaphors in 2D and provide links to papers and studies on the metaphor. Also the envelopes have descriptions on them. Most of them lead to systems that use a very different spatial concept than the city itself. She enters one of the envelopes and finds herself floating in space with molecule-like structures floating around her.

She doesn’t understanding the concept behind these structures, and finds not documentation in the room either. She activates the magic window on the business card again to ask the room owner. (Figure 3) A chat connection is opened. As she states her question the room owner pulls out a document from a pouch she is carrying and hands it through the magic window mentioning that this is not linked it yet. The document shows the abstract of a paper on the metaphor our user is using with a link to the full article. Intrigued by the abstract the user activates the link and teleports to the full paper.

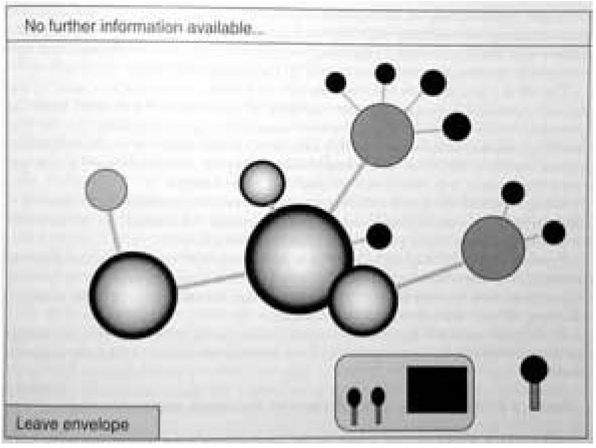

Figure 3: Inside containers information might be presented using very different metaphors. In this envelope, the user sees information represented as a stick-and-ball molecule model. Note the button to leave the envelope (lower left) and the magic window on the lower right (see scenario).

6.3. Exploring the city

Instead of choosing teleportation in the last section our user could have used the subway metaphor which would give her better feedback on the distance covered. Using the subway the user realizes she is traveling to a very far destination in the city. Upon arrival she skims the article for a while and then decides to look around this neighborhood a little. The article is located in an area with personal buildings. The topics here are very diverse. She summons a tramway car to get a tour. The car moves through the neighborhood on a path determined both by the most commonly used paths and her interest profile. Very soon she discovers a buildings owned by big names in her field. She decides to get off the tour and approaches the buildings, but finds that some of them are surrounded by a fence or have locked entrances signaling that entrance is not possible without invitation. She drops her virtual business card into a mailbox with a request to gain access. As the owner is not currently available she has to come back later. Strolling around she finds a sidewalk cafe that seems to be a popular hangout for this neighborhood. She enters and soon is engaged in a free-wheeling chat discussion on computers and society.

7. Implementing an Information City

The Information City is a conceptual metaphor. It is unlikely we will ever see a complete implementation of all ideas and structures described so far. However there are several possibilities how to implement parts of the Information City and to use the ideas of the city for information systems. In this section we outline possibilities and problems of an implementation using either a textual or a graphical virtual environment.

7.1. Using a text-based virtual environment

Textual virtual environments evolved out of text-adventure games. In these game environments the player (user) is represented by a *player character* or *avatar* moving through an environment described by text. In modern games of this kind many players use the environment at the same time, and can encounter and interact with each other.
In this paper we use MOO (MUD Object Oriented) as a general term for multi-user textual virtual environments. MOO systems can be used for almost real-time communication and distributed work. They easily cope with a large number of users and they commonly develop into *virtual communities* [13] and *social places* [20].

The MOO user controls the *player character* by issuing commands in a more or less natural command language. Also the outcome of commands like "look at book" are described textually as if the player character actually had done something [7, 12, 13]. This seemingly old-fashioned command-line interface provides access to a rich and detailed virtual world. A MOO is a low-bandwidth multi-user environment of a size and complexity that is still unrealizable in graphical environments.

Using a modern client software it is possible to associate non-textual information to objects, locations and activities in the MOO environment. An example for such a system is the *Juggler* system developed by one of the authors [17, 18].

Combined systems like Juggler allow us to experiment with navigational metaphors in a very flexible way while the actual content is provided in a separate window. A text-only representation of a spatial metaphor is flexible enough to realize aspects of the metaphor that may be very difficult to realize graphically.

For example while most user will find a textual representation more difficult to get used to than a well-designed graphical representation it can be more efficient. For example a description like

```
You are in a hidden chamber in the great pyramid. Gold glitters wherever you look.
```

will evoke a strong mental image in the user although the description is so brief. A graphical representation of the same room will have to be very elaborate to evoke the same image in the user. This difference is one of the major advantages of the textual environment. The two major applications for this advantage are the representation of read wear and the enactment of magic features.

Special exits in a MOO system commonly are described as portals or teleporters and the transition is described to be like the transporter beam in the Star Trek movies. Special exits in a MOO seldom give an indication which direction the user is headed in. However, users commonly associate an upward movement and long distance with teleportation as described in [58]. Typical examples for transitions are

```
Merlin gets more and more translucent and with a soft *plot* he is gone.
```

The downside of a textual representation of magic features is that it does not easily provide ongoing feedback. On the other hand one of the general problems of the read wear concept is that it is difficult to realize convincingly in a graphical form. Slight changes in a visualization over time may well go undetected and it is unclear how to best represent aging of information. A graphical representation of age actually can mislead users as was found in [53] because a simulation of aging of paper by giving a document a yellowish tone caused users to perceive older documents as more important than newer ones. Read wear is closely related to the concept of aging information and thus it is also difficult to find a good graphical representation of read wear. However, it is relatively easy to represent read wear textually.
In the Juggler system [17] room exits (the connections between distinct locations) remembered how often they had been used. Frequently used exits were pointed out in the room description which made major navigational paths through the environment visible like in the following example.

```
The exits {east} and {west} seem to be used above average.
```

Descriptions of this kind hint at navigation patterns in the user population. These patterns (perceived as often traveled paths) help users to orient themselves and also make it easier to locate rooms of general interest.

Similarly the Juggler system used read wear on bulletin boards. Bulletin boards could be read either in the textual environment or though a Web gateway. The Web gateway used a graphical representation of the read wear, whereas the MOO client used a textual representation. From interviews with the system's users we know that users were unable to make sense of the graphical representation without further explanation. The textual read wear on the other hand was understood by everybody without training [18].

As these examples show there are some advantages in using a text-based system to implement (parts of) the Information City. Especially for magic features and read wear, two concepts both of which are difficult to realize convincingly in a graphical version, the textual representation provides an easy to realize and use alternative. This is a desirable feature especially for experimental implementations. For a mass-market implementation of the Information City text-only implementation probably will not be successful though.

7.2. Using a graphical virtual environment

Not too long ago a realization of the Information City even on a high-end computer would have been material for a science fiction movie. With the ongoing boom of the World-Wide Web, the acceptance of VRML (Virtual Reality Modeling Language) and the availability of very powerful desktop machines a purely graphical realization of an Information City is within reach.

The technical possibility does not alleviate us from addressing those problems we can avoid in a text-only environment: the design of magic features and read wear. We believe that especially the design of believable magic features requires transferring knowledge from the fields of cartoon animation and (movie) special effects. Interface designers will have to cooperate with specialists outside the traditional user interface fields, a strategy that proved to be helpful already in the past. For examples and source materials see [10, 11, 55, 56].

Information contained in an Information City often will be in the form of text. A graphical realization therefore has to make sure this text is presented in a readable form to the user. Text on an information wall has the disadvantage that users will see if at an oblique angle, which will bring back the problems of early spatial metaphors trying to extend the flat desktop (see section 3.3).

The problems of text on flat surfaces in three-dimensional space have been addressed in several publications, like the World Processor [59], in a Web browser that is used within a virtual environment [4] and in the Web book interface [9], just to name a few examples. The Web book uses the perhaps obvious solution that the user first selects an object to read (a book showing a Web page) and that object is moved so it hangs in space perpendicular to the user's eyes.

7.3. Building a city

An important unresolved issue for a complex structure such as the Information City is how to actually create the city structure. It is unlikely that a complete Information City can be created in one step. Instead, like in a real city, it has to develop over time out of the interaction of its users (inhabitants?) with the environment.

Initially the city might be only a small collection of basic infrastructure residing in empty space (the *void*). Large districts do not have a fixed location but are arranged as floating islands separated by a small stretch of undefined space. This prevents them from ever growing into one monolithic city. Districts therefore can grow without being restricted by adjacent districts. The connection between these islands is provided through flying or the subway.

An important consideration is minimal infrastructure that ideally should be available in every container -- at least in a simplified form. It consists of at least one object that can serve as landmark. The district, and every other container thus develops out of a kind of *seeding kernel*.

An alternative to this scheme is to provide a well-developed organization of a district from the very start, like many American cities are built according to a similar basic layout. Without a concept like the void such a structure may get crowded in certain places and will lead to the formation of satellite districts that are separated from each other although they may contain related information.

Yet another idea for the creation of a city (suggested by Thomas Erickson) is to provide a 'deserted city', modeled after a real city. Users adapt this structure as needed when they populate the space, changing the city's character over time.

These ideas still leave many issues unresolved. For example it is not clear how to determine how much influence a single user can have on changes in the environment and how fast changes are allowed to happen in the environment. We assume that many of these problems will get tackled by social protocols that evolve out of the interactions in a user population, very much like a democracy eventually formed in the MediaMOO system. These issues are beyond the scope of this paper.

8. Conclusions

We described a conceptual spatial metaphor, called the Information City. This metaphor is based on knowledge transferred from the fields of architecture and city planning and designed to support navigation in the resulting virtual city environment. Contrary to other city metaphors in the literature we define a detailed ontology of city elements, describing how each element pertains to a navigational structure. Navigation tools provide the infrastructure to move in the Information City.

The city is supposed to scale up well because we define magic features that provide short-cuts through the city space. We also define a structural magic feature, called an envelope, that can contain a space that is larger than itself. The envelope and the void, stretches of undefined space, allow the city to expand and shrink without becoming unrecognizable and therefore hard to navigate. We also give suggestions how the city initially can be set up.

A set of both stationary and movable information providers is defined to fill the city structure with information. We also describe the read wear concept which visualizes usage information in the environment.

The Information City is designed to be a multi-user environment. It is not a sterile information graveyard but a social space. The information rich environment of the Information City supports users in giving directions, in recognizing landmarks and so forth.

We also discuss several problems implementors will have to address when using either a textual or a graphical virtual environment. In particular these issues are the design of magic features, the visualization of usage information (read wear) and text readability, that is the design of the interaction with the environment and of the representation of certain processes and information. As with all information systems also an Information City -- although it is modeled after real life cities -- will require good design in its details to be usable. With the rapid technological progress a fully graphical implementation of the Information City is probably possible today but without a well designed implementation of magic features and read wear an Information City will fall short of its potential.

Even with our quite detailed description of the Information City there are many design decisions left open to the implementor. Among them is the design of movable information objects and the behavior of transportation metaphors that are not in use. Another interesting design issue is the meaning and development of linear elements in the city and how they grow into a certain direction over time.

In section 7 we compared advantages and disadvantages of textual and graphical implementations of spatial metaphors, in particular the Information City. Questions like the design of movable information objects and the meaning of linear elements are typical example of problems where a textual virtual environment can be used to experiment with a metaphor before creating a full-fledged graphical version for the mass-market. We believe that textual virtual environments provide us with a tool for simple and rapid testing of spatial metaphors in a multi-user environment.

Acknowledgments

Parts of the work for the Information City has been funded by a research grant from the Austrian "Fonds zur Förderung der wissenschaftlichen Forschung" (Dr. Erwin Schrödinger Stipendium), Grant J01021-MAT. The groundwork for this publication was laid in the first author's Ph.D. thesis, advised by Peter Fleissner and Andrew Frank of the Vienna University of Technology. Without their continued support this work would have been impossible. It was further stimulated and supported in many discussions with Thomas Erickson and Werner Kuhn.

References

1. Alexander, C. A City is Not a Tree, in Kaplan, S. and Kaplan, R. (Eds.). Humanscape - Environments for People. Ulrich's Books Inc., Ann Arbor, 1982, pp. 377-402.
2. Andrews, K. Visualizing Cyberspace: Information Visualization in the Harmony Internet Browser, InfoVis'95, IEEE Press, Atlanta, 1995, pp. 97-104.
3. Andrews, K., Pichler, M., and Wolf, P. Towards Rich Information Landscapes for Visualising Structured Web Spaces, Proc. of 2nd IEEE Symposium on Information Visualization (InfoVis'96), San Francisco, CA, 1996, pp. 62-63.
4. Angus, I.G. and Sowizral, H.A. VRMosaic: Web Access from within a Virtual Environment, InfoVis'95, IEEE Press, Atlanta, 1995, pp. 59-64.
5. Barreau, D. and Nardi, B.A. Finding and Reminding - File Organization from the Desktop, SigCHI Bulletin, 27, 3, (1995), pp. 39-43.
6. Bernstein, M. Enactment in Information Farming, Hypertext'93, Seattle, 1993, pp. 242-249.
7. Bruckman, A. and Resnick, M. Virtual Professional Community: Results from the MediaMOO Project, Convergence, 1, 1, (1995) .
8. Card, S.K., Robertson, G.G., and Mackinley, J.D. The Information Visualizer: An Information Workspace, CHI'91, 1991, pp. 181-188.
9. Card, S.K., Robertson, G.G., and York, W. The WebBook and the Web Forager: An Information Workspace for the World-Wide Web, CHI'96, ACM Press, Vancouver, BC, 1996, pp. 111-117.
10. Chang, B.-W. and Ungar, D. Animation: From Cartoons to the User Interface, UIST'93, Atlanta, 1993, pp. 45-56.
11. Clanton, C. and Young, E. Film Craft in User Interface Design -- CHI'94 Tutorial Notes, Boston, 4/94, 1994 .
12. Curtis, P. Mudding: Social Phenomena in Text-Based Virtual Realities, 1992 .

13. Curtis, P. and Nichols, D.A. MUDs Grow Up: Social Virtual Reality in the Real World, Xerox Parc, 1993 .
14. Dieberger, A. Navigation in Textual Virtual Environments using a City Metaphor, PhD Thesis, Vienna University of Technology, 1994.
15. Dieberger, A. On Magic Features in (Spatial) Metaphors, SigLink Newsletter, 4, 3, (1995), pp. 8-10.
16. Dieberger, A. Providing Spatial Navigation for the World Wide Web, in Frank, A.U. and Kuhn, W. (Eds.). Spatial Information Theory - Proceedings of COSIT'95. LNCS 988, Springer, Semmering, Austria, 1995, pp. 93-106.
17. Dieberger, A. Browsing the WWW by interacting with a textual virtual environment - A framework for experimenting with navigational metaphors, Proc. Hypertext'96, Washington DC, 1996, pp. 170-179.
18. Dieberger, A. Supporting Social Navigation on the World Wide Web, International Journal of Human-Computer Studies, 46, (1997), accessible at http://www.cc.gatech.edu/gvu/reports/techreports97.html, pp. 805-825.
19. Dieberger, A. and Bolter, J.D. On The Design of Hyper"Spaces", Communications of the ACM, 38, 8, (1995), p. 98.
20. Erickson, T. From Interface to Interplace: The Spatial Environment as a Medium for Interaction, COSIT'93, Elba, 1993, pp. 391-405.
21. Furnas, G.W. Effective View Navigation, CHI'97, ACM Press, Atlanta, GA, 1997, pp. 367-374.
22. Henderson, D.A. and Card, S. Rooms: The Use of Multiple Virtual Workspaces to reduce Space Contention in a Window-Based Graphical User Interface, ACM Transactions on Graphics, 5, 3, (1986), pp. 211-243.
23. Hendley, R.J., et. al. Narcissus: Visualizing Information, InfoVis'95, IEEE Press, Atlanta, 1995, pp. 90-96.
24. Hill, W.C. and Hollan, J.D. Edit Wear and Read Wear, CHI'92, ACM Press, Monterey, 1992, pp. 3-9.
25. Houde, S. and Solomon, G. Working Towards Rich & Flexible File Representations, InterCHI'93, ACM Press, Amsterdam, 1993, pp. 9-10.
26. Ingram, R. and Benford, S. Improving the Legibility of Virtual Environments, Proc. of the 2nd Eurographics Workshop on Virtual Environments, Monte Carlo, 1995.
27. Ingram, R. and Benford, S. Legibility Enhancement for Information Visualization, Proc. IEEE Visualization'95, Atlanta, GA, 1995.
28. Ingram, R., Bowers, J., and Benford, S. Building Virtual Cities: Applying Urban Planning Principles to the Design of Virtual Environments, Proc. VRST'96, ACM Press, HongKong, 1996.
29. Jerding, D.F. and Stasko, J.T. The Information Mural: A Technique for Displaying and Navigating Large Information Spaces, InfoVis'95, IEEE Press, Atlanta, 1995, pp. 43-50.
30. Johnson, J. How Faithfully Should The Electronic Office Simulate The Real One?, SigCHI Bulletin, 19, 2, (1987), pp. 21-25.
31. Jones, W. and Dumais, S. The Spatial Metaphor for User Interfaces: Experimental Tests of Reference by Location versus Name, ACM Transactions on Office Information Systems, 4, 1, (1986), pp. 42-63.
32. Kuhn, W. and Blumenthal, B. Spatialization: Spatial Metaphors for User InterfacesCHI'96 Course Notes, 1996 .
33. Kuipers, B. Modeling Spatial Knowledge, Cognitive Science, 2, (1978), pp. 129-153.
34. Lakoff, G. and Johnson, M. Metaphors we live by, Univ. of Chicago Press, 1980.
35. Laurel, B. Computers as Theatre, Addison-Wesley, 1991.
36. Lynch, K. Good City Form, MIT Press, 1981.
37. Lynch, K. The Image of the City, MIT Press, 1982.
38. Mackinley, J.D., Robertson, G.G., and Card, S.K. The Perspective Wall: Detail and Context Smoothly Integrated, CHI'91, 1991, pp. 173-179.
39. Malone, T. How Do people Organize Their Desks? Implications for the Design of Office Information Systems, ACM Transactions on Office Information Systems, 1, 1, (1983), pp. 99-112.
40. Maltz, D. and Ehrlich, K. Pointing The Way: Active Collaborative Filtering, CHI'95, ACM Press, Denver, CO, 1995, pp. 202-209.
41. Maurer, H. HyperWave: The Next Generation Web Solution, Addison Wesley, 1996.
42. Mountford, S.J. and Gaver, W.W. Talking and Listening to Computers, in Laurel, B. (Ed.) The Art of Human Computer Interface Design. Addison-Wesley, 1990, pp. 319-334.
43. Mukherjea, S. and Foley, J.D. Visualizing the World-Wide Web with the Navigational View Builder, GVU report, GVU-TR 95-09, 1995, accessible at ftp://ftp.gvu.gatech.edu/pub/gvu/tech-reports/95-09.ps.Z.
44. Nielsen, J. Usability Engineering, Academic Press, Boston, 1993.
45. Nielsen, J. Multimedia and Hypertext - The Internet and Beyond, Academic Press, Cambridge, 1995.
46. Nievergelt, J. and Weydert, J. Sites, Modes, and Trails: Telling the User of an Interactive System where he is, what he can do, and how to get to places., in al., G.e. (Ed.) Methodology of Interaction. North-Holland, 1980, pp. 327-338.
47. Poblete, F., Chignell, M.H., and Chaffey, A.V. No Free Ride: When Information Visualization Doesn't Promote Learning of Information Structure, University of Toronto, Dept. of Industrial Engineering report, TR #96-01, 1996 .

48. Resnick, P., et. al. GroupLens: An Open Architecture for Collaborative Filtering of NetNews, CSCW'94, Addison Wesley, Chapel Hill, NC, 1994, pp. 175-186.
49. Resnick, P. and Varian, H.R. Recommender Systems, Communications of the ACM, 40, 3, (1997), pp. 56-58.
50. Robertson, G.G. and Mackinlay, J.D. The Document Lens, UIST'93, Atlanta, 1993, pp. 101-108.
51. Rose, D.E., Borenstein, J.J., and Tiene, K. MessageWorld: A new Approach to Facilitating Asynchronous Group Communication, CIKM'95 (Conf. on Information & Knowledge Management), Baltimore, MD, 1995, pp. 266-273.
52. Smith, D.C., et. al. Designing the Star User Interface, in Baecker, R.M. and Buxton, W.A.S. (Eds.). Readings in Human-Computer Interaction. Morgan Kaufmann, 1987, pp. 653-661.
53. Solomon, G. New Use for Color, in Laurel, B. (Ed.) The Art of Human Computer Interface Design. Addison Wesley, 1990, pp. 169-178.
54. Tauscher, L. and Greenberg, S. Revisitation Patterns in World Wide Web Navigation, CHI'97, ACM Press, Atlanta, GA, 1997, pp. 399-406.
55. Thomas, F. and Johnston, O. Disney Animation: The Illusion of Life, Abbeville Press, New York, 1981.
56. Tognazzini, B. Principles, Techniques, and Ethics of Stage Magic and Their Application to Human Interface Design, InterCHI'93, ACM Press, Amsterdam, 1993, pp. 355-362.
57. Tognazzini, B. Tog on Software Design, Addison-Wesley, 1996.
58. Tromp, J.G. and Dieberger, A. MUDs as text-based spatial user interfaces and research tools, Journal of Intelligent Systems, 5, 2-4, (1995), pp. 179-202.
59. Verlinden, J.C., Bolter, J.D., and Mast, C.v.d. The World Processor - An Interface for Textual Display and Manipulation in Virtual Reality, TU Delft report, 93-55, 1993 .
60. Wexelblat, A. and Maes, P. Footprints: Visualizing Histories for Web Browsing, MIT Media Lab, Software Agents Group, 1997,.