A Pattern of Islands: Exploring Public Information Space in a Private Vehicle

John A. Waterworth
Department of Informatics, Umeå University,
S-901 87 UMEÅ, Sweden
E-mail: jwworth@informatik.umu.se

Abstract. Increasingly, we are entangled during our daily working lives in a web of distributed, networked information sources offering varied means of electronic communication. We have become information explorers in this complex, electronic world-at-large. In addressing the problems of navigation and orientation that these developments raise, a user interaction model is developed which also deals with some of the issues of workspace management addressed by systems such as Rooms and the Information Visualizer. In this model, the world-at-large is represented as Information Islands, each of which contains Buildings which themselves house items of information. The user explores this world in a Vehicle which has two views of the world; the public, comprehensive view, and a private, customised view. The Vehicle can also be seen as the user's own private workspace. Wherever the user wanders in cyberspace, he is always at home.

1 Introduction

The last few years have witnessed a fusion of research on Virtual Reality, Multimedia and Hypermedia. Historically, hypertext triggered interest in information exploration, by allowing the user to follow cross-referential links between and within bodies of text. This trend, towards non-linear and unpredictable routes through collections of information, led to a vision of the user as a wanderer in a world wide information space known as 'cyberspace', In this chapter I refer to this newly-evolving space as the informational 'world-at-large' or just 'the world'.

Although there are still many media-related issues to be resolved [Hardman et al., 1994; Waterworth, 1992], the burgeoning development of multimedia technology means that it is becoming realistic to think of the items in information space as having presence in a variety of media: sound, video, animation, text, and so on. The world of information that an individual may explore extends across a variety of information providers, systems and regions of the globe; and it is growing rapidly.

There are many ways in which the informational world-at-large may be represented to users, and some suggestions and studies (mostly the former) of how

well users fare with different representations. Smith and Wilson [Smith & Wilson, 1993], for example, discuss a range of representations from schematic, two dimensional overviews to three dimensional spatial walkthroughs. They suggest that what constitutes a successful representation depends on the nature of the database and of the users. This chapter describes a particular users' model of the informational world, Information Islands.

The two principal aspects of the world-at-large that affect the success of user interactions within that world are the way the information is structured and the way the structure is represented to users. The work described in this chapter assumes that the world-at-large can be structured in a particular way; a way that provides a good match to a particular form of representation - as archipelagos of islands, each of which is populated with Buildings that may be many storeys in height. The justification for this assumption, and the limitations this imposes on the generality of the approach described, are discussed in Section 5 of the chapter.

The 'Information Islands' model is designed to address questions of how we might make this increasingly complex and diverse informational world accessible and intelligible, while at the same time providing an individual user with a comfortable, customisable environment. A major aim of the model was to provide a flexible, rich and expandable paradigm for the organisation and presentation of information and services; an approach that could be applied at various levels in a compatible way and could convey direct means of locating and using services and applications.

Several approaches to the spatial representation of abstract information have been developed, but there has been little or no systematic usability testing. One of the earliest and best known is the SemNet system [Fairchild et al., 1988]. The other widely known work in the field is the Perspective Wall, Cone Trees and the Information Visualizer, all from the same team at Xerox PARC [Card et al., 1991; Robertson et al., 1991, 1993]. The Workshop on Spatial User Interface Metaphors in Hypermedia at the 1994 European Conference on Hypermedia Technology provides a useful summary of some more recent approaches [Dieberger, 1994].

In the remainder of this chapter I describe aspects of the Information Islands model in some detail. The basic concepts of Archipelagos, Information Islands, Buildings, and Vehicles are covered in the next two sections.

2 Archipelagos, Islands and Buildings

The world-at-large is seen as a set of Archipelagos, each composed of Information Islands. Each Archipelago represents a set of broadly related entities, providing a clear, top-level classification of what is available in the world-at-large and where it is to be found - an overall orientation that is easily accessible to both the novice and the experienced user. Each major class of service or application exists as an Archipelago. Examples might be Entertainments, Government Services, Information Services, Communications, Medical, and Financial Services.

Archipelagos are collections of Information Islands. The size of an Archipelago depends on the number (and size) of the Islands of which it is composed.

Figure 1 shows a highly schematic representation of the top level of the world. When implemented, this will be represented as a pseudo-3D world through which the user can navigate (see later figures).

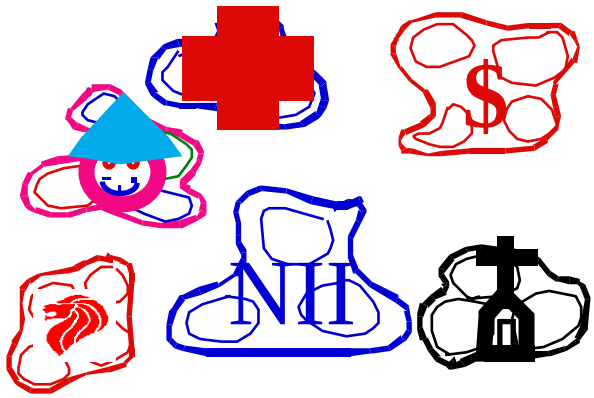

Figure 1 - Top level of the world - 6 Archipelagos

Each Island generally contains only one subclass of service. Islands are the fundamental semantic unit of the world-at-large. Users will become familiar with this world mostly by learning the location of Islands with the kinds of services they use or are interested in. Each Island contains one or more Buildings. Some Islands may be representations of the services offered by particular providers - Provider Islands. An example might be a particular information provider's Island located near other Information Services Islands.

Each Building contains a set of information sources or services related to a particular topic or application focus. Examples might be Weather Building, Sports Building, Stocks and Shares Building. Buildings on a particular Island will have distinctive appearance (shape, colour, graphics, text). All Buildings have common features including a Store Directory and an Information Counter (see Figure 2). The Store Directory allows users to browse and select from what is available in a Building. The Information Counter is a public agent that searches for information in response to requests from users (cf. 'Information City': [Waterworth, 1992; Dieberger, 1993; Dieberger and Tromp, 1993]. Buildings will contain standard features to assist in navigation and item location (cf. [Musil and Pigel, 1993]).

Entities (Archipelagos, Islands, Buildings) get physically bigger the more items they contain. At the top level, major application areas are shown as Archipelagos, i.e. collections of Islands. Each Archipelago represents a major topic or application focus. Each Archipelago is formed by placing a boundary around the Islands from which it

is composed. Each Archipelago has a distinctive colour that provides a context and reminder to the user of the focus he has chosen.

As the user zooms in for more detail, the view of Archipelagos is replaced by a view of the Islands from which the selected or central Archipelago is composed. Intermediary views provide realism and orientation as the user zooms down. In the intermediate view, Islands appear as “raisins” in the Archipelago “fruit cake” (see Figure 1).

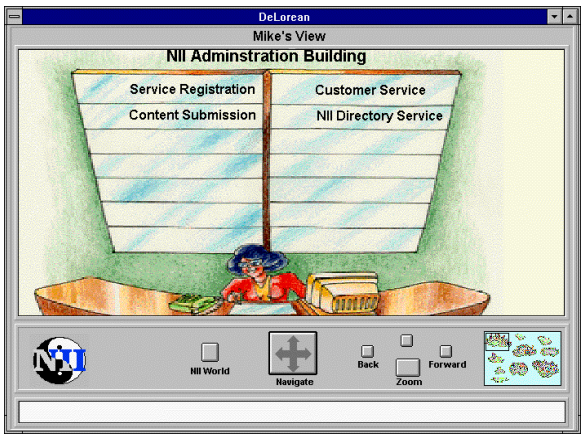

Figure 2 - A Store Directory and Information Counter

When a single Archipelago is shown, the Islands from which it is composed are represented separately. A view of a single Island is a map of the collections of services provided and which are represented as Buildings (see Figure 3).

Buildings that are related are clustered together into no more than ten villages. Each Building contains no more than 20 Floors, and each Floor generally contains a set of related services.

The user views the available services by zooming down and selecting a particular Building, which will be of a different colour from that of its neighbours. He enters the foyer (the background will retain the colour of the Building it belongs to) and can then either browse the Store Directory or consult the public search agent at the Information Desk. The Store Directory presents a list of the service types available on each Floor of the Building.

At each Floor there is a ‘lift lobby’ where users consult a Floor Directory (like the Store Directory, but listing individual services), to invoke the service they require.

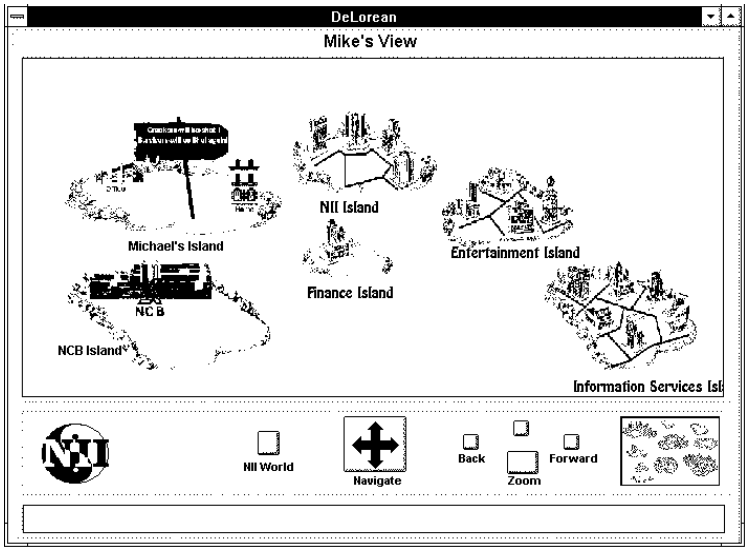

Figure 3 - Islands and Buildings

3 Vehicles

An important part of interacting with the world-at-large is the idea of exploration, selection and collection of items of interest to the individual user. These items may be services, information or particular configurations of applications. One common way of catering for this need for a private selection from a public world is to demarcate part of the world as private, and allow the user to collect items and configure that private area. This is one of the essential ideas behind the Rooms concept [Henderson and Card, 1986]. However, such an approach is limiting. Users must navigate to their own area frequently, bringing back items they want to collect, then venture out again into the world-at-large. Here the disadvantages of a spatial metaphor can outweigh the advantages: because the users' private space is part of the global information space, they frequently have to move around to switch between their own perspective and the higher levels of organisation. Use (which always involves a user) is being confounded with organisational level (which includes a User level). Use should be possible at any level, at any time. A private area at a particular location in the world is not the best way of supporting individual customisation, because it does not provide enough support for exploration of, and collection from, the fast-expanding, networked, multimedia world-at-large.

To overcome these problems, the concept of private Vehicles was developed; these can be thought of as transparent, mobile, personal workspaces (but with a difference!). Vehicles combine the idea of a private collection of information and configuration of services (customised workspace) with that of multi-level

navigational device and customised information viewer. The user is always in his (or her) Vehicle, and therefore always has access to both public and private worlds. Items can be transferred between these worlds without navigating space. A key aspect of the model is that the user has a filtered way of looking at the same spatially-arranged world that occupies public space. It is a manipulable viewpoint rather than a specific place (cf. [Nagel, 1986]).

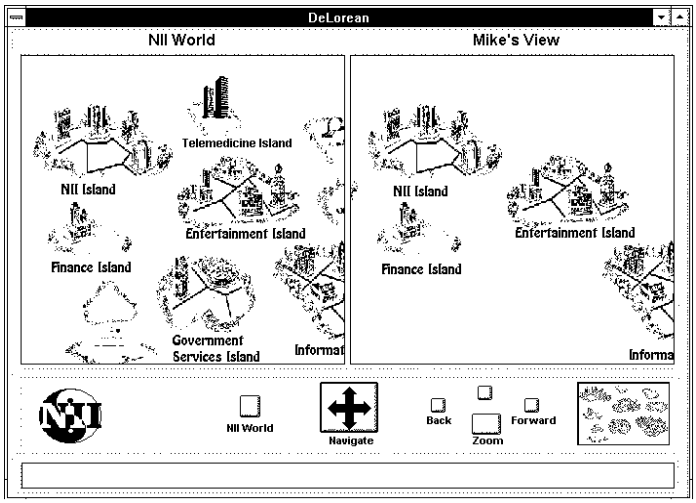

Figure 4 - Inside the User's Vehicle: Public and Private Views of Islands and Buildings.

The user in his Vehicle has two Views of the world outside - a public "God's Eye" View that contains everything that is available, and a personal View showing only those items that the user has selected as of interest or use (see Figure 4). He has only one set of navigation and viewing controls; the user chooses upon which View or Views they act. Although there are two Views, there is only one world. The private View and the God's eye View are different perspectives on the same world; the former is filtered and limited, the latter is a complete display at the level of detail on which it is focused. The user can choose to have a split screen showing both Views simultaneously, or alternate between the two.

The user can 'yoke' the two Views together so that the public View and the private View are both from the same viewpoint[1], changing together as the user navigates or inspects materials at different levels. This can be useful when he wants to know what else is available at a place, other than the things he has already chosen to include in his View. This is also useful during customisation, when he can fly around the world-at-large and select things that he will then see included in his own View. At other times, he will select one or other of the two Views to be updated as his Vehicle moves, but not both. The View that is selected (private or public) will be the one that

[1]Viewpoint is, literally, the position in virtual space from which Views are taken.

is affected by the navigation controls, the other will remain focused on where it was when last selected. He can use the public View as a navigational overview while exploring in detail with the private View. Alternatively, he can have his private View as an overview and move around the world via the public View collecting items to add to his private world. Selecting ‘yoke’ will cause the less-recently-selected View to be updated to match that of the more-recently-selected View. This means that navigation can be done on either View, and the other View aligned to that perspective when required.

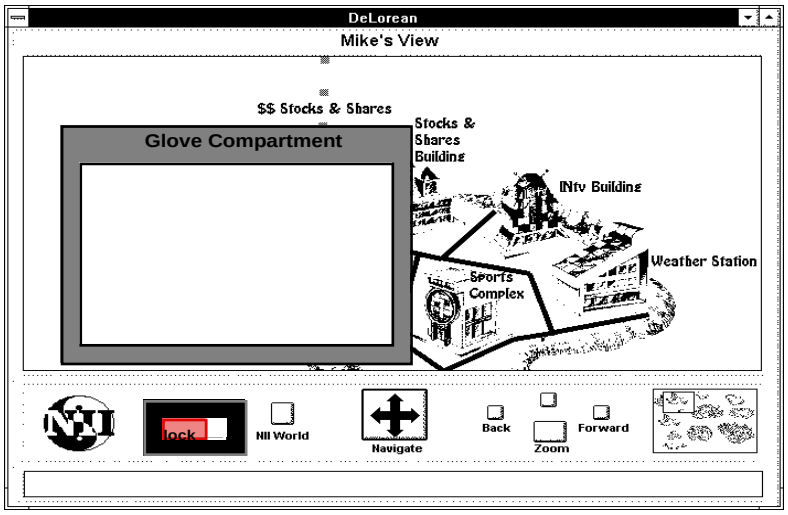

Figure 5 - Glove Compartment (Open and Empty).

The provision of both a private and comprehensive public View means that the user has access to a customised world, as well as the world-at-large. This customised world is a subset of the world-at-large, selected by the user but retaining the layout and grouping inherited from the larger world. Apart from this simplified View, the user may want instant access to a few frequently used services and applications. Two mechanisms are provided for this: the Vehicle’s Memory and the Glove Compartment (see Figure 5). The Vehicle’s Memory is a list of places the user wants the Vehicle to remember, so that they can be rapidly revisited without the need for navigation. When he is at a particular location, at whatever level in the hierarchy, he may select the ‘Memorise’ option, resulting in that location being added to the memory list. At a later time, he simple clicks on that item in the list to instantly move from wherever he is to that location. A simple ‘Forget’ option allows locations to be removed from the list.

In addition to the Vehicle’s Memory there may be particular applications or services to which the user wishes to have instantaneous access, and/or which may be used at a variety of locations. Such items can be stored in the Vehicle itself and so are

always with the user wherever he may be in the world-at-large. The Glove Compartment is located at the lower left of the display, to the left side of the navigation controls. When closed, the 'Open' option is display. Selecting this causes a moveable window to appear, displaying the contents of the Glove Compartment. Applications are stored in the Glove Compartment by the user dragging their icon from a navigation window onto either the open window or the Glove Compartment feature on the dashboard. Items are removed by dragging out of the open window and dropping anywhere else.

4 Information Exploration

Information exploration (which includes locating services and applications, as well as tracking down more specific items of information) may be divided into two kinds of activity: navigation by the user, and search by the system. These correspond to the two main styles of human-computer interaction: direct manipulation (navigation) and conversation (to elicit and receive the results of search by the system). [Waterworth and Chignell, 1989; 1991] expand on these two kinds of activity and point to ways in which they can be satisfactorily combined.

Under the Information Islands model, users navigate around the world-at-large in Vehicles. Public search agents are located in the world-at-large, private agents are found only in users' Vehicles.

4.1 Navigation by the User: Time and Space, not just Unrelated Places

A fundamental part of the Islands model is the idea of the user navigating in his Vehicle - finding his own way to the things in which he is interested. His Vehicle is his means of navigation and it contains controls that permit three kinds of navigation: moving around the world, viewing different levels of the world, and travelling in history. The Islands model uses spatial location as a means of representing the classification of items within the world-at-large. Users will frequently move around this spatial representation to locate items of interest - from Island to Island, from Building to Building, and within Buildings. Navigation is thus an important feature of interacting with the world. This type of travel implies the need for controls (in the user's Vehicle) that permit smooth movement around the Islands environment in the horizontal plane - move north, move south, move west, move east.

Travel of this kind conveys to the user a compelling experience of information space, not just jumping from place to place. The standard experience of browsing networked information sources (such as the World Wide Web), and indeed of following links to destinations in any large hypermedia system, is of arriving somewhere in information space, without any spatial relationship to other locations. In other words, users "teleport" instantly to another location, but may be disoriented about how one location and its contents relate to another and to the information found

there. A truly spatial informational model, such as Information Islands, takes care of the rhetorics of arrival. Instant travel to known locations (see Section 3) is merely a convenient extra (magic) feature.

As well as moving around the world, users will also want to view the world at different levels of detail - they may want to see the world as a whole and then focus on one Island of interest, then one Building, then one Floor, and so on. Having been involved at a detailed level, they may then want to see more of the surrounding items, perhaps to locate related items. Or they may want to revert to a high-level overview before returning to a more detailed level. This second type of travel requires smooth movement in the vertical plane - move up and move down. Additionally, users will want to move instantly down to an item of interest, and may also require the ability to move instantly up for a high level view.

When users have moved from place to place, corresponding to their engagement in different activities, they will want to retrace their steps without having to themselves locate where they were last (and before that, and before that, and so on...). Having gone back in time one or more steps, which means revisiting previous locations, they may then want to move forwards in time to resume previous activities - retracing steps to the locations visited before. Travelling in time (more accurately: travelling back and forth along their own interaction history) requires only two controls: move back in time and move forward in time.

For the three types of navigation described above, two basic controls are needed which, in keeping with the metaphor, can be seen as Vehicle navigation devices. They are located on screen and selected by mouse. Additionally, they could be mapped onto key clusters.

Each control has four direction buttons in a cross-shaped grouping (see Figure 5). Move North, Move South, Move East and Move West correspond to the four buttons of the horizontal movement control. At the corresponding locations of the other control, north and south positions are used for Move Up (the hierarchy) and Move Down (the hierarchy), respectively. West is used for Move Back (in time) and East for Move Forward (in time). Instant Move Down is achieved by mouse clicks on the chosen location. Double-clicking on the Move Up button signifies Instant Move Up (to a high level view).

4.2 Public and Private Agents

Public agents are the same for all users, whereas private agents can be configured to meet the needs and behaviour of the particular user they serve. In browsing and otherwise navigating the world, users employ direct manipulation to find the items they want. When agents are employed, the user engages in something analogous to a conversation, to specify his requirements which the agent then attempts to satisfy. The operation of agents is fairly simple, since the range of queries handled and responses provided are limited.

There are two types of public agent: the Tour Agent and the Information Counter Assistant. The Tour Agent resides in “space”, outside the world of Archipelagos and Islands. It is only present at the top-most level of viewing the world. Its function is to provide general information about the layout of the world-at-large. The user can request information about what is available and where it is located. The Tour Agent will arrange one of various “tours” to meet the user’s needs (see [Waterworth, 1992] p.179). The Information Counter Assistant is similar to the Tour Agent but the topics on which it seeks to provide information are limited to the contents of the Building in which it is located.

Private agents are found in the user’s Vehicle. The user can consult an agent to aid navigation or to invoke a service or item of information directly by having an agent search for it. Additionally, they can monitor changes and update the user when requested or at regular intervals (say, daily or weekly). The user can specify the type of changes about which he wishes to be informed and can modify this specification depending on how satisfied he is with the results (see [Fischer and Stevens, 1991] for an early example).

5 One World, Many Possible Representations

In presenting any visualisation of complex information that is at all intelligible to users we need two things: a way (or ways) of structuring the information, and a way (or ways) of representing such structure. How a particular way of structuring information is presented to the user comprises the world (or users’) model. With the Information Islands model, the information is structured into a hierarchy presented as Archipelagos, Islands, Buildings, and Floors. But the information could be structured differently, and different world models used to present that structure. Additionally, the world model may be filtered to provide various Views.

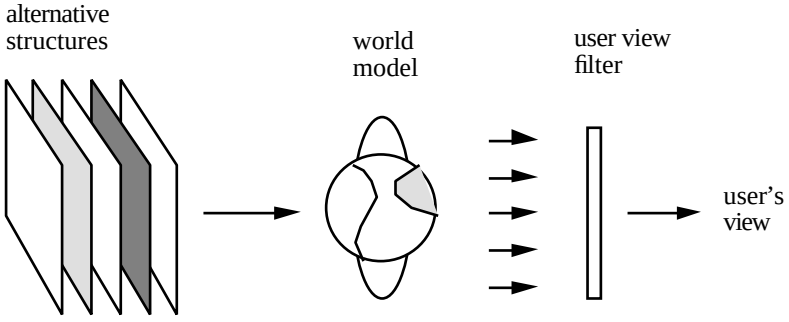

Figure 6 - Structures, World Model, User's View

With the Islands model, the interface is adapted to particular users’ needs by selection of appropriate Views, with the organisational structure and user world

model remaining fixed and common to all users. In theory, however, selection could actually be made from alternative realisations at all three levels. So a user might choose to have a set of data organised according to one or more possible structures (or the same structure with different attributes, such as geographical data, by turn-over, by product type, for example), might then select from a range of world models (sets of nested boxes, tree graph, transparent buildings) by which that structure would be represented and, finally, would choose suitable Views for the task in hand. Figure 6 illustrates these three levels of information presentation diagrammatically.

There are several unanswered questions arising from this work. Is a single hierarchical structure realistic? What are the advantages of Information Islands versus other world models? Would forests, trees, and leaves have been any different (e.g. the "Dataforest"; [Rifas, 1994]). Would more than two Views give additional benefits?

If we assume a hierarchy of ten Archipelagos, with twenty Islands per Archipelago, twenty Buildings per Island and twenty Floors per Building, we have the necessary scope for a large number of individual information items to be located in the world. With twenty items per Floor, we have 1.6 million items. Relaxing the restrictions on items per Floor by having sub-sets of items accessed by two submenus after the initial Floor Directory selection, and expanding the world to a maximum of 20 Archipelagos, would allow us to accommodate over a billion individual items. Can users navigate in such a world?

But this view of expansion is unrealistic. The Information Islands model was designed to meet a particular need. It was assumed that the world would start life relatively empty and would then gradually expand, as providers offered information and other services. In this sense, it is rather like a plan for a city. But it is not clear to what extent development will match the original planning. As Alexander has pointed out, "A city is not a tree" - not a simple hierarchy that grows according to predictable rules. As providers offer services, and users gravitate to the things they are interested in (i.e. willing to spend time and/or money on), the original plans are likely to be heavily modified by market pressures. Like a pleasant city, the world-at-large should evolve to meet the needs of users and providers alike. But not all cities are pleasant, and the balance between central planning and market-led evolution is not easy to strike. An additional vital question is whether users can actually find their way around such an evolved model.

6 Conclusions

A major problem in designing a personal interface to the complex, networked, multimedia electronic world that is now widely available is how to show what is available and how it is organised. The Information Islands model represents the cyberspatial world-at-large as archipelagos of islands. Because each class of service is shown as a separate Island with distinctive features, we expect users quickly to acquire an understanding of the range of services available. By flying around the

Islands they obtain an overview of (one possible organisation of) the world-at-large. They are encouraged to explore and, in so doing, should soon become familiar with the main features.

One advantage of the Islands metaphor is that it supports the idea of a loose collection of related entities (Islands) each having similar basic properties. These basic properties make all Islands predictable from experience of only one. The classification hierarchy naturally follows: the world is a group of Archipelagos, each Archipelago is composed of a number of Islands, each Island contains a group of Buildings, each Building contains different information or services. By exploring Buildings within the context of Islands, users should soon develop a sense of the hierarchical organisation of the world. This type of geographical metaphor has the advantage of being highly intuitive and easily understandable, as compared to more abstract or limited approaches such as graphical charts, systems of folders, or traditional sets of directories (see [Waterworth and Chignell, 1989]). However, it is not clear how users will cope with the model as it evolves through use.

Another issue addressed by the model is that of how users customise their own workspaces so that selected items, such as the services or applications they want to use frequently, are directly available to them. The customary approach to this problem is to give users a private area which they can customise to meet their own needs. But if this area is represented as one location in the spatial world shown as the interface, they must do a great deal of navigating between this private world and the world-at-large. Users need to be able to alternate freely between these two perspectives, so that they can avoid unnecessary complexity while at the same time having easy access to the public world when they want it. Previous approaches tend to enclose the user within his private area so that he loses sight of the broader range of services available. When he is exploring the world at large the process of collecting items and placing them in a particular location can be laborious.

The Information Islands model solves this problem by providing a private as well as a public View of the world, and through the concept of a Vehicle. Each user has his own Vehicle with which he explores the world of Information Islands and Buildings. As he navigates around the world, the public View shows all the items in existence. In addition to this “God’s Eye” View, the user has a filtered, private View of the world, showing only those items of interest to him. Items from the public View are simply selected for inclusion in the private View; they do not have to be transported back to a particular, private location (although limited storage space is provided in the Vehicle itself). The user is always in his Vehicle. Not only is he a wanderer in information space, he is a nomad of no fixed abode.

"We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time."

T S Eliot, *'Little Gidding'*.

Acknowledgements

Thanks to Gurminder Singh for creative brainstorming sessions, to Jayne Waterworth for useful discussions of the model, and to Erik Stolterman for valuable comments on an earlier version of this chapter.

References

[Card, Robertson & Mackinlay, 1991] Card, S K, Robertson, G G and Mackinlay, J D (1991) The Information Visualizer, an Information Workspace. In *Proceedings of CHI'91* Conference on Human Factors in Computer Systems (New Orleans, April 1991), 181-188. New York: ACM.

[Dieberger, 1993] Dieberger, A (1993) The Information City - a step towards merging of hypertext and virtual reality. Poster at *Hypertext '93*.

[Dieberger, 1994] Dieberger, A (1994) Report of the Workshop on Spatial Metaphors at *ECHT'94 - the European Conference on Hypermedia Technology*, September 1994, Edinburgh, UK.. Accessible through the World Wide Web only (http://www.gatech.edu/lcc/idt/Faculty/andreas_dieberger/Workshop.ECHT94.html).

[Dieberger & Tromp, 1993] Dieberger, A and Tromp, J G (1993) The Information City Project - a virtual reality user interface for navigation in information spaces. In *Proceedings of the Symposium Virtual Reality Vienna*, December 1-3, 1993.

[Fairchild, Poltrock & Furnas, 1988] Fairchild, K M, Poltrock, S E and Furnas, G W (1988) SemNet: Three-Dimensional Graphics Representations of Large Knowledge Bases. In *Cognitive Science and its Applications for Human-Computer Interaction*, R. Guindon (ed). Hillsdale NJ: Lawrence Erlbaum Associates.

[Fischer & Stevens, 1991] Fischer, G and Stevens, C (1991) Information Access in Complex, Poorly Structured Information Spaces. *Proceedings of ACM CHI'91* Conference, 63-70, New York: ACM.

[Hardman,, Bulterman & van Rossum, 1994] Hardman, L, Bulterman, D C A and van Rossum, G (1994) The Amsterdam Hypermedia Model: Adding Time and Context to the Dexter Model. *Communications of the ACM*, 37 (2), 50-63.

[Henderson & Card, 1986] Henderson, D A and Card, S K (1986) Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-Based Graphical User Interface. *ACM Transactions on Graphics*, 5 (3), 211-243.

[Musil, S & Pigel, 1993] Musil, S and Pigel G (1993) Virgets: Elements for Building 3-D User Interfaces. In *Proceedings of the Symposium Virtual Reality Vienna*, December 1-3, 1993. Also available as TR 93/13, Vienna User Interface Group, Lenaugasse 2/8, A-1080 Vienna.

[Nagel, 1986] Nagel, T (1986) *The View from Nowhere*. New York: Oxford University Press.

[Rifas, 1994] Rifas, L (1994) The Dataforest: tree forms as information display graphics. In Dieberger, 1994.

[Robertson, MacKinlay & Card, 1991] Robertson, G G, MacKinlay, J D and Card, S K (1991) Cone Trees: Animated 3D Visualizations of Hierarchical Information, *Proceedings of CHI''91* (New Orleans, Louisiana, 28 April - 2 May, 1991), ACM, New York, pp. 189-194.

[Robertson, Card & Mackinlay, 1993] Robertson, G G, Card, S G and Mackinlay, J D (1993) Information visualization using 3D interactive animation. *Communications of the ACM*, 36(4):56-71, April 1993.

[Smith & Wilson, 1993] Smith, P A and Wilson, J R (1993) Navigation in hypertext through virtual environments. *Applied Ergonomics*, 24 (4), 1993, 271-278.

[Waterworth, 1992] Waterworth, J A (1992) *Multimedia Interaction: Human Factors Aspects.* Chichester, UK: Ellis Horwood, 1992.

[Waterworth, 1995] Waterworth, J A (1995) Viewing Others and Others' Views: Presence and Concealment in Shared Hyperspace. *Workshop on Social Contexts of Hypermedia*, 16-17 February 1995, Umeå University, Sweden.

[Waterworth & Chignell, 1989] Waterworth, J A and Chignell, M H (1989) A Manifesto for Hypermedia Usability Research. *Hypermedia*, 1 (3), 1989, 205- 234.

[Waterworth & Chignell,1991] Waterworth, J A and Chignell, M H (1991) A Model of Information Exploration. *Hypermedia*, 3, 1991, 35-58.

