

Foreword

The Human Touch: Reflections on i^3

The Machine-centred Mind Set

At the Chicago World Fair of 1933, the official motto was: “Science Finds – Industry Applies – Man Conforms”. To many of us today this seems quite shocking, yet it has been the driving force of much development in the last century.

In particular, if you look at the rise of computing over the last 50 years, you will see that, on the whole, development has been extraordinary, but fairly straightforward: it can be characterised as trying to make “faster and faster machines fit into smaller and smaller boxes”.

Starting from the time of the ENIAC, one of the colossal computers of the 1940s, most IT progress has been driven from the point of view of the machine. Since then things have changed – but perhaps not really that much. Even if computers can today calculate many times over what was possible a few years ago, and the machines have become somewhat less obtrusive, much of the “mind set” has stayed the same. It is the visions of huge calculating machines spanning massive rooms, trying to recreate an absolute artificial intelligence, that still haunt much of the thinking of today.

Clearly, it is difficult to shake off old mind sets.

Alternatives

Alternatives to the idea of fitting computing into ever smaller boxes can mainly be attributed to Mark Weiser. In his paper, “The Computer for the 21st Century”, he outlined notions of how computing could become integrated into the fabric of everyday life by becoming completely distributed into the environment. In this way computing would become “ubiquitous”. More recently, similarly inspired work on “tangible media”, by Hiroshi Ishii has emerged from the MIT Media Lab. Apart from this, the technological revolution of GSM and the mobile phone has also had its share of making information technology come out of its “traditional shell”.

Alternatives to the machine-centred view to computing were also starting around the same time, such as the “anthropocentric” ideas proposed by Mike

Dertouzos at the MIT Computer Science Lab; and in a similar vein, cognitive scientist, Don Norman has been pointing out the lack of well-designed information environments.

Roughly at the same time, but from a different perspective, we started thinking about how to give technology more of a “human touch”. Now, in principle, this should not be that difficult, as technology is after all, made by humans. In practice, however, one has to go quite far to break down the machine-centred and box-centred ways of thinking.

We decided that the only way to attack the problem with any significance was to try to invert the picture completely – that is, to start thinking from the human point of view and work outwards. Our idea of “human centredness” was that it should nurture technological innovation but within a broader context of human values and aspirations. This was not the same as “user” driven, or “defined by user needs”, all of which tend to become stuck in improving the status quo, but not growing beyond it. At the same time, we also wanted to make sure to break out of the box-centred ways of thinking as much as possible and avoid doing “traditional HCI”, which was mainly involved in improving computers as they were.

Our ideas were designed to balance questions of technically “how”, with questions of “why?” and “what for?”. And the aim was to see if we could start restoring the balance between people’s inventiveness to make new machines, with the essence of being human. Our questions became rather: How can we reach a better and more fulfilling balance between technology and people? What could be new ways of thinking about the problems? What could be the new paradigms that could lay the paths for further research and development?


The i³ Research Programme

It is along these lines that we launched our first call for proposals back in 1995. Our general aim was to look at the relationship between people and information technology in the future: how could people access and use information, and exchange things with others using information technology as a medium?

A clear break was needed to get out of stale thinking. Therefore, we called for new *paradigms* of interaction and research on new interfaces between people and the world of information. We also asked how such work could intertwine human, societal and technological elements into one dynamic research activity. One of the main quotes from our call for proposals was:

The goal of i³ is to research and develop new human-centred interfaces for interacting with information, aimed at the broad population (1996).

To help define a specific research agenda, we first had a competition for more specific visions of the future. “Connected Community” and “Inhabited Information Spaces” were selected as the two visionary themes on which we based a subsequent call for research projects. Even though it took some time



to have an extra layer of calls for proposals, in retrospect it was better to “reculer pour mieux sauter”.

The two selected themes had similar yet contrasting underlying philosophies. The Connected Community theme, proposed by a team headed by Irene Mac William (Philips Design, Eindhoven) and Marco Susani (Domus Academy, Milan) asked: forget about virtual environments and trying to fit people into some artificial world – how can we help people in their everyday environment, and integrate technology into this? The idea is to understand how information and communication tools start making a difference when they are embedded in a real context, and start being more meaningful for actual people and communities. How can technology enhance these environments and activities, rather than replace them?

The other schema, Inhabited Information, proposed by a team headed by Tom Rodden (University of Nottingham) took a slightly different perspective. It stated: the Internet and the Web already represent a suspended reality, and people want to participate more in these spaces. Given that this is a reality, how could it evolve in the future? How could we make it more accessible to the broadest possible public, and make it socially interactive for large groups of people, in meaningful ways? And in similar spirit to the first theme, how can such environments link to the physical everyday world rather than be removed from it?

At a later stage, we decided to supplement the research with an emphasis on learning. We wanted to explore new relationships between learning and technology. The idea was that a lot could be learnt about designing new interfaces by looking at how children interact, play and learn. Similar ideas had been experimented in a Lego context by Seymore Papert of the MIT Media Lab. In 1997, we decided to have a call on experimental school environments (ese). This centred around learning for very young children, in fact, the 4–8-year-old age range. This age range struck us as being particularly challenging because at this stage children don’t have too many of the adult preconceptions of the world, and are still open to new things. Young children have a different kind of “language” – a form of communication and expression from which adults can learn a lot. From this we wanted to gain insights about how to design meaningful interaction tools for the population at large. The header of our call was:

The aim of i³-ese, is to research new kinds of IT-based tools designed to enable new approaches to learning, focussing on the age range of 4 to 8 (1997).

From each of these programmes we selected a number of individual research projects. Together these spanned many universities, research centres and companies across Europe, and involved a mix of people from many walks of life – artists, designers, computer scientists, game companies, technology companies, experimental schools, teachers and children, people in communities, etc. At the same time all these different outlooks were united by the common vision: of exploring new relationships between people and technology.

Grains of the Future

In this book you will find some examples of work in particular from the Inhabited Information Spaces Grouping. It is interesting to see how some of these ideas are still “futuristic” and others have started to become part of mainstream thinking and made their way into products.

Some people say that you can find “grains of the future” in the present today – the only problem is, where do you start to look? One of the potential advantages of this book is that by looking at the research developments stretching out into recent past, one can identify how some grains developed into trends of the present, and other are still just emerging.

For those still interested in seeking out “grains of the future”, this book will be a valuable source.

Jakub Wejchert
Information Society DG
European Commission

Jakub Wejchert grew up in Ireland, with a family background of artists and architects, of Polish origin. He studied natural science at Trinity College Dublin, specialising in physics, and holds a doctorate (modelling of non-linear networks) from the same institution. Later he worked in the USA with IBM research, working on computer graphics and interface design. He joined the European Commission in 1992. At the Future and Emerging Technologies unit, he set up and managed a number of research programmes such as *i³* – intelligent information interfaces; *i³* – experimental school environments; and the “disappearing computer”. He now works as an advisor on vision and strategy to one of the Directors in the Information Society Programme. Jakub lives in Waterloo, south of Brussels, with his wife and three sons.

The opinions expressed here are those of the author and do not necessarily reflect the position of the European Commission.

WWW3D and the Web Planetarium

Mårten Stenius and David Snowdon

2.1 Introduction

This chapter will describe an Inhabited Information Space based around a 3D visualisation of a portion of the WWW. The system was originally called WWW3D (Snowdon et al., 1996), and evolved from an experiment in immersive 3D web browsing into a richer inhabited information space called the Web Planetarium during the i³ eSCAPE project (Stenius et al., 1998). We include it here as an example of an inhabited space supporting a common modern activity – web browsing – and also as an example of how a basic visualisation was extended to correct deficiencies in the earlier system.

WWW3D started life as an experiment in 3D web browsing. We were thinking about some of the original ideas behind VRML (Virtual Reality Modelling Language, www.vrml.org), namely that it would be a sort of 3D web. In fact many people working on VRML were inspired by the Cyberspace detailed in William Gibson's novels:

Cyberspace. A consensual hallucination experienced daily by billions of legitimate operators, in every nation, by children being taught mathematical concepts . . . A graphic representation of data abstracted from the banks of every computer in the human system. Unthinkable complexity. Lines of light ranged in the nonspace of the mind, clusters and constellations of data. Like city lights, receding (Gibson, 1986).

However, VRML became more a means to represent 3D models with hyperlinks between them, than a richer way of experiencing the web. It so happened that as we were thinking about this during the summer of 1996 we were asked to provide a demonstration for SICS's (the Swedish Institute of Computer Science's) new large-screen immersive VR system. This system, while having impressive graphical performance, did not support interaction via a standard 2D GUI (Graphical User Interface); we therefore decided to create a web browser that was entirely 3D. While

the result was not a web browser as usable as Netscape's or Microsoft's for standard web browsing, the prototype had several interesting features and evolved into a more capable system during the i³ eSCAPE project. In this chapter we will describe the original system, WWW3D, and its evolution into the Web Planetarium.

Just as a normal web browser, WWW3D allowed users to follow links and view the web pages associated with those links. However, rather than simply showing the current web page the user was exploring, WWW3D also visualised the structure of the portion of the web that the users had explored, historical information showing when pages had been last viewed and (thanks to the DIVE environment for which WWW3D was written) other users browsing the web at the same time. Note that WWW3D was never intended to visualise the web – the web is far too large for this to be meaningful, instead we simply visualise the portion of the web that users have explored and are exploring. We assume that existing search engines are sufficient to locate new pages of interest. The two key features of WWW3D were its method for representing web pages in 3D and the method for organising multiple web pages in 3D space. We will describe how this was done before continuing with a description of how the basic prototype was improved.

2.2 Producing a 3D Representation of a Web Page

Many HTML tags provide semantic information about the marked-up text – that is, what the text *is* rather than how it should be displayed on the page. 2D web browsers will typically use this information to determine what style (i.e. font, font style, font size, etc.) to use when drawing a particular piece of text.

WWW3D uses the information contained in HTML tags to produce a representation of the document in 3D space. A web document is represented as a sphere that is labelled with the document's title. The contents of the document is placed around the inside surface of the sphere. Displaying large amounts of text in a satisfactory way is difficult in current VR systems, so textual information is currently represented by icons that can be unfolded to reveal the entire text. The first few words of the piece of text are displayed under the icon to give some indication of the contents. Images are displayed by texture mapping them onto polygons on the inside surface of the sphere. Finally, links to other documents are represented as icons labelled with their destination.

To reduce the visual complexity of the virtual environment WWW3D makes extensive use of level of detail (LOD) operations. When viewed from outside, a document is represented as an opaque sphere and the actual document contents are not displayed. When a user enters a document to view it, the sphere is drawn in wire-frame so that the rest of the virtual environment is still visible.




Figure 2.1 A WWW3D representation of a web document seen from the inside. The spheres in the background represent other web pages.

Figure 2.1 shows the contents of a web document as displayed by WWW3D. Figure 2.2 shows the display generated when a user reads some of the text comprising a web document.

2.3 Browsing the Web Using WWW3D

When a user selects a link icon, WWW3D creates a new sphere representing the target document and places it near the document from which the user selected the link. In order to indicate the structure of the portion of the Web that the users have explored, WWW3D draws arrows between the spheres representing linked documents. If the documents are resident on the same Web server then the arrow is drawn in blue, otherwise

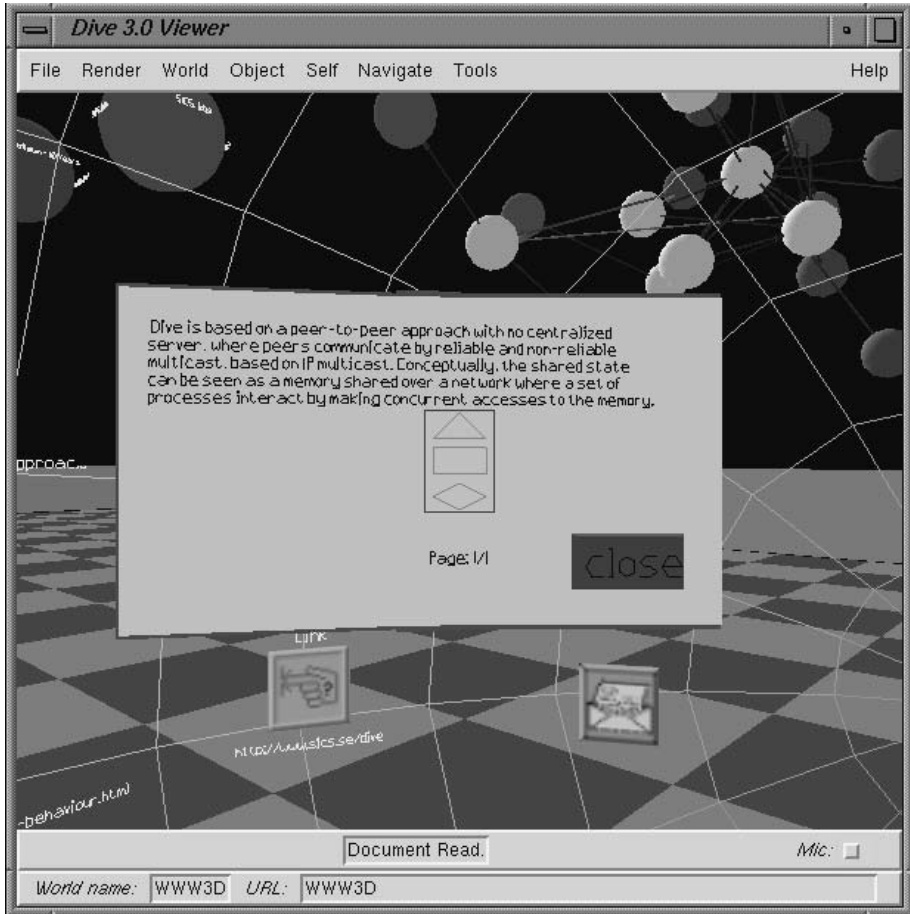


Figure 2.2 Reading part of the text of a web document.

it is drawn in green, thereby helping to provide additional information on the structure of the documents that the user has explored. In addition to this, the brightness of the arrow is dependent on the time since a user last followed that link thereby providing users with a visual representation of their browsing history. If WWW3D fails to fetch a document then a small red arrow is attached to the source document to represent the “broken” link .

WWW3D is implemented using multiple lightweight threads so users do not have to wait for a document to be retrieved before selecting another link. This behaviour is essential if multiple users are to be able to browse independently. Users are also free to navigate through the space and browse other documents while waiting for a document to be retrieved.

As WWW3D parses a newly retrieved document, it checks for links to documents that the user has already explored and draws arrows to represent them. This means that at any given moment the complete set of links between documents is displayed without users having to follow every link. This is intended to aid the user by indicating links between documents that the user might have been unaware of. This also has the result that several users can be browsing different parts of the web and yet any links between the sets of documents they are exploring will be displayed. This might be useful since users will then have a visual representation of possible common interests.

To produce an acceptable layout of the set of linked documents an incremental version of the Force Directed Placement (FDP) (Fruchterman and Reingold, 1991) algorithm is used. Links between documents act like spring forces that result in linked documents being moved closer together. Documents exert repulsive forces on one another which prevents documents being placed closer together than a user-specified minimum separation.

At regular (user-specified) intervals WWW3D applies the FDP algorithm to refine the inter-document layout for a specified number of iterations. The more links a document has, the greater “inertia” it is considered to have when the FDP algorithm is applied. This has the result that heavily referenced documents are less likely to move and provides some stability to the visualisation. In addition damping is applied to try and prevent large changes to the visualisation for a given iteration of the FDP algorithm – this helps to prevent the user from becoming disorientated. The result of this is that the inter-document layout gradually evolves over time to produce clusters of inter-linked documents. Instead of running the FDP algorithm until it converges (which could take a long time) only a specified number of iterations are executed. The rationale for this is that the space will change anyway as users browse new web pages so it is not necessary to force users to wait until the layout stabilises before allowing them to continue browsing. If there is no change to the space (no new web pages) then the system will gradually converge to a stable state. However, the FDP algorithm suffers from the disadvantage that it can take a great many iterations to converge to a stable state.

Figure 2.3 shows the display produced by WWW3D after the user has browsed a number of documents. The FDP algorithm has resulted in the formation of clusters of closely linked documents. The colours of the documents provide some indication of how long ago the user last visited them.

Between invocations WWW3D stores information on the current set of documents, the links between them and the current 3D layout to a file. When a new instance of WWW3D is started, it reads this history file and displays the structure as it appeared in the last session. The contents of the documents found from the history file are not retrieved until the user

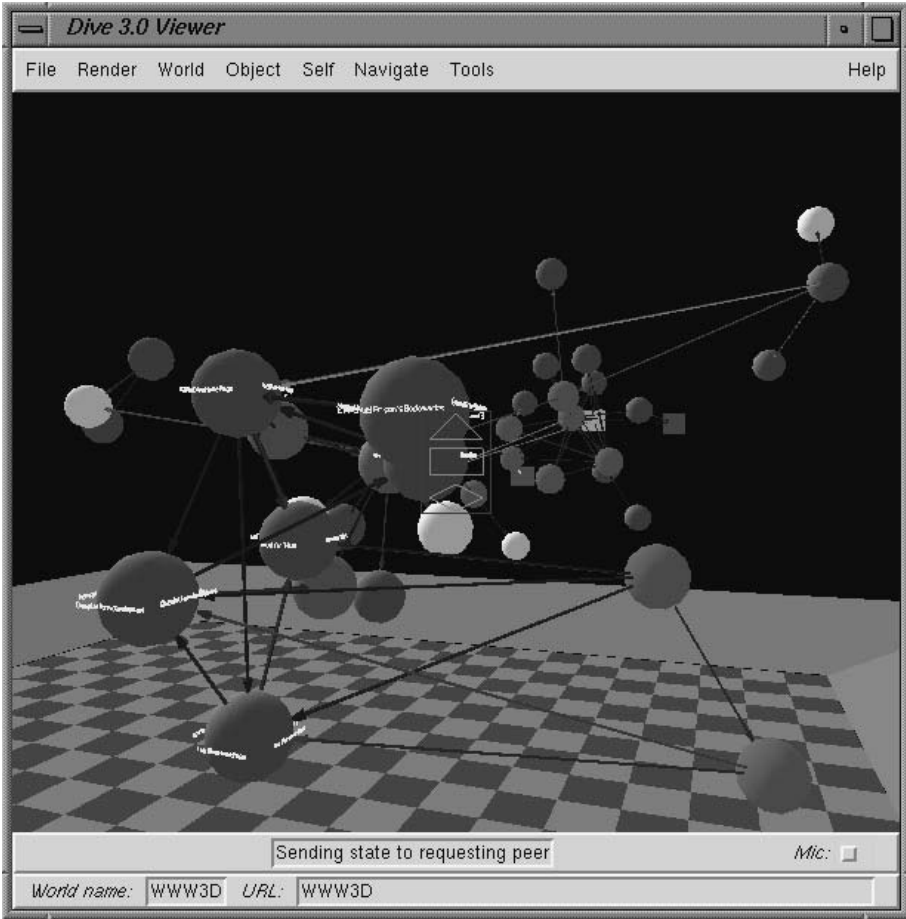


Figure 2.3 A collection of web documents, showing the links drawn between them.

enters a particular document. Since the document representations are opaque this process is invisible to the user except for an occasional delay in seeing a document's contents.

2.4 Improving Scalability

One problem with the WWW3D prototype is that it suffers from scalability problems both in terms of display and number of users. Since a single DIVE application is responsible for management of the visualisation, it is vulnerable to becoming overloaded if several users simultaneously place demands on it. In addition to this the complexity of the display increases as documents are added, lessening the advantage of

the displayed web structure information and making it harder to browse. Even though WWW3D makes extensive use of level-of-detail (LOD) operations for individual nodes, there will still be a point where the world becomes too complex to be rendered on even the most powerful hardware. For this reason we created a new version of WWW3D running on the MASSIVE-2 CVE (Benford et al., 1997a) that attempted to solve the display problem by making more extensive use of LOD effects combined with clustering of web pages.


We decided to extend the use of LOD beyond the contents of a single web page to encompass groups of web pages (e.g. all pages on the same server, same domain name, owned by the same user etc.). Doing this meant that instead of showing all the pages in a website, or all the personal pages (home page and related pages) for a user, we could display an object that represented the whole group. The user would only see the contents of the group when a certain criteria was met – in this case getting close enough to the group. This meant that when a user was outside a cluster their client had no need to know anything about the web pages in the cluster or the users that were currently browsing web pages in the cluster. This would have the following advantages:

- A reduction in the visual complexity of the world.
- A consistent metaphor since clusters looked and behaved similarly to individual web pages when users entered them.
- Increased interactivity since the computational load on the client machine was reduced.
- Reduced network bandwidth since MASSIVE-2 assigned separate multicast communication channels to each cluster. This meant that a user who could not see the interior of a cluster had no need to receive network updates concerning objects or users contained in that cluster.

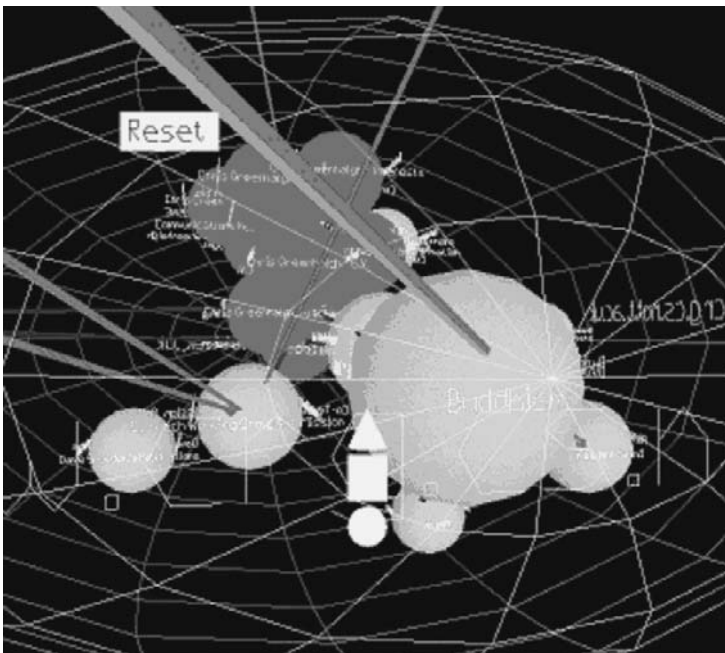
We used a simple scheme based on URLs to perform our clustering. A more advanced alternative would be to explicitly consider the *legibility* of the CVE and cluster accordingly as is done by LEADS (Ingram and Benford, 1995) which is also capable of adding additional objects such as landmarks to aid users navigating through a CVE.

Our new implementation was based on two fundamental building blocks, containers and *links*. Containers have the ability to enclose other objects in a spatial sub-region of the virtual space. Links connect the containers. A special form of container is the *page*, which represents a web page and its contents. A *cluster* is a *container* that is able to treat its contents as a sub-visualisation complete with layout, other containers and links. Figure 2.4 shows a cluster both from a distance and as the user gets sufficiently close to see the contents.

In order to simplify the representation of links between clusters, we introduced *compound links*. A *compound* (or aggregate) *link* between two clusters is formed whenever there exists one or more links between the



a



b

Figure 2.4 A cluster abstraction dissolves to display its contents when the user approaches: (a) the exterior of the cluster; and (b) the view seen by a user as they get closer to the cluster.




Figure 2.5 A compound link (AB) between two clusters (A , B).

contained objects of the two clusters. Thus, a compound link both serves as a link structure between two clusters and as an abstraction for one or more lower-level links between the contents of these clusters, this is illustrated in Figure 2.5. Figure 2.6 shows a collection of linked clusters.

When refined, the current application structure approach should lend itself to parallelisation: letting different processes handle different clusters. This in turn would enable a distribution of calculations on different hosts to allow more complex visualisations.




Figure 2.6 A set of interconnected cluster abstractions. Their sizes roughly reflect the number of objects (here pages) contained within.

2.5 The Web Planetarium: Creating a Richer Visualisation

The Web Planetarium extends the WWW3D concept with the idea of the system becoming a gateway between different virtual spaces. The metaphor of a planetarium was used to create a new, more informative and more visually appealing look and feel, since the system was also going to be used in public settings using a large, dome-like display (Schwabe and Stenius, 2000). Nodes in the Web Planetarium display can represent either web pages, 3D models (in VRML or DIVE format) or 3D DIVE worlds that can be jumped to. The Web Planetarium can therefore also act as a space of portals to other spaces and provide a means of navigation between different virtual worlds. Compared to WWW3D the Web Planetarium provides a more visually interesting virtual space and different layout options.

2.5.1 Visual Differentiation of Nodes

One problem with the WWW3D visualisation is that all nodes look alike. The only distinguishing feature is colour and that colour represents the time since the last access and not a feature of the node itself. When creating the Web Planetarium, it was felt more important for users to be able to quickly scan the space for potentially interesting pages than to show the usage information. The Web Planetarium replaces the use of colour by texture maps extracted from the web pages. Alternatively, the sphere representing a node can be replaced by a 3D model. Not only does this result in a more interesting environment but it allows the users to distinguish more easily one node from another and gain some idea of what the node represents. Figure 2.7 shows an example Web Planetarium view in which the texture maps clearly distinguish nodes from one another.

A simple algorithm is used to select the image to texture onto a sphere – the first image found in a web page is used as the texture. In the case of a user's home page this image is typically a photo of the user – this is easily visible in Figure 2.8 in which several of the spheres clearly show people's faces.

Obviously, a simple algorithm such as this will sometimes fail to produce an interesting image for one of several reasons.

- There is no image on the web page.
- A banner advert is selected instead of an image representative of the content of the page.
- The image selected is a piece of web decoration (a line, a bullet point) instead of something really representative of the subject of the web page.




Figure 2.7 The 3D layout alternative emphasises the concept of a planetarium.

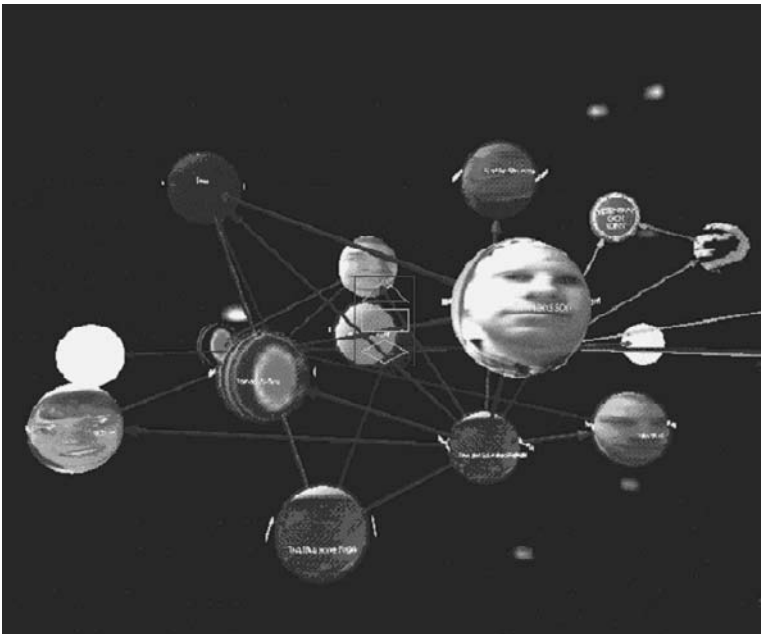


Figure 2.8 A screenshot of the Web Planetarium. Compared to the original WWW3D, one of the obvious differences is the use of texture mapping to give more information about the contents of web pages instead of a simple colour indicating the time since last access.

It would be hard, probably impossible, to produce an algorithm guaranteed to work, but there are a number of simple filtering operations that could be used to ensure that in many cases unsuitable images are rejected.

- Reject images that are too small (according to predefined parameters) or are long and thin. This extremely simple approach was used successfully in the CWall electronic notice board (Snowdon and Grasso, 2002) and rejected most web decoration .
- Use a blacklist of URLs and URL components to reject banner advertisements or configure a proxy server that performs this function (Hood, 2000). There are websites that maintain such blacklists (@Man, 2000) so each user would not have to shoulder this burden.
- Reject images that do not come from the same server as the web page. This is a very simple approach that will work in many cases, but it may sometimes reject legitimate images.

An alternative approach could be to use a conventional web browser – embedded within the application or running in parallel – to generate a snapshot of the web page, including background colour, text layout, images and so on. The snapshot would then be used as an icon (or “thumbnail”) to be pasted on the outside of the 3D page representation. This method would have the advantage of creating a stronger connection to the original design of the web page as a whole.

2.5.2 The Web as a Road Network

WWW3D and the Web Planetarium can generate very tangled and convoluted visualisations of nodes. Although users can manually reposition nodes this is not a solution that can be applied to large spaces. The Web Planetarium offers users the options of restricting the FDP layout algorithm to a horizontal plane, thereby giving the layout shown in Figure 2.9. A disadvantage of this approach is that in certain cases the FDP algorithm may take longer to converge.

2.5.3 Hybrid Browsing

The original aim of WWW3D was to provide an acceptable and novel web-browsing experience using an entirely 3D interface. This made an interesting demo and made sense on a large screen or head-set based system but is less useful on the more common desktop-VR systems in which the VR system is just another window on the user’s 2D desktop. The Web Planetarium therefore allows users to combine the 3D view of the web structure with a traditional web browser view of the contents of a web page. This is illustrated in Figure 2.10 which shows a Web Planetarium view and a web browser (Netscape) view of the same web page side by side.




Figure 2.9 The 2D layout alternative results in a road-like network of sites and links.

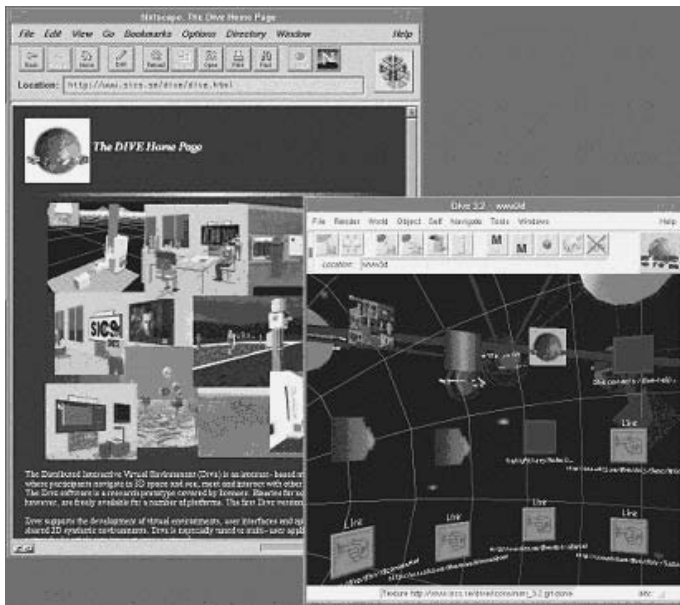


Figure 2.10 Using a conventional web browser as a side-viewer for HTML documents encountered in the Web Planetarium. Whereas the 3D view of the HTML document (to the right) gives an overview of the content and interaction points for links, Netscape provides a detailed close-up of the current page (to the left).

2.6 Conclusion

This chapter has described WWW3D, a simple 3D web browser and its evolution into the Web Planetarium. Starting from a humble beginning WWW3D has been changed to address the issues of scalability and producing a more pleasing and informative visualisation. In addition it can now be seen as providing a means for navigating between disjoint 3D worlds as well as between web pages.

However, the current version of the Web Planetarium is not without problems. One major issue is related to the viewing of 2D standard web content. Neither the current 3D view nor the separate browser and 3D views side by side really provide a convenient browsing experience. There is still work to be done to find the optimal way to provide information about web structure and a view on web page contents in an integrated way.

The second problem with the Web Planetarium is that while it is perfectly possible for multiple users to share a single instance of the Web Planetarium there is no support given to the interrelation of different planetaria nor to different views on the same planetarium. The use of subjective visualisations (Jää-Aro and Snowdon, 2001) may provide a technique for merging planetaria since the same web page could be presented differently to different users, thus preserving individual layouts while allowing a degree of interaction between users.