

**CENTRE FOR ADVANCED
SPATIAL ANALYSIS
Working Paper Series**



Paper 1

**MODELLING
VIRTUAL URBAN
ENVIRONMENTS**

**Michael Batty
Martin Dodge
Simon Doyle
Andy Smith**



Centre for Advanced Spatial Analysis
University College London
1-19 Torrington Place
Gower Street
London WC1E 6BT

Tel: +44 (0) 171 391 1782
Fax: +44 (0) 171 813 2843
Email: casa@ucl.ac.uk
<http://www.casa.ucl.ac.uk>

<http://www.casa.ucl.ac.uk/modelvue.pdf>

Date: January 1998

ISSN: 1467-1298

© Copyright CASA, UCL.

ABSTRACT

In this paper, we explore the way in which virtual reality (VR) systems are being broadened to encompass a wide array of virtual worlds, many of which have immediate applicability to understanding urban issues through geocomputation. We sketch distinctions between immersive, semi-immersive and remote environments in which single and multiple users interact in a variety of ways. We show how such environments might be modelled in terms of ways of navigating within, processes of decision-making which link users to one another, analytic functions that users have to make sense of the environment, and functions through which users can manipulate, change, or design their world. We illustrate these ideas using four exemplars that we have under construction: a multi-user internet GIS for London with extensive links to 3-d, video, text and related media, an exploration of optimal retail location using a semi-immersive visualisation in which experts can explore such problems, a virtual urban world in which remote users as avatars can manipulate urban designs, and an approach to simulating such virtual worlds through morphological modelling based on the digital record of the entire decision-making process through which such worlds are built.

1. Definitions

Virtual environments are digital simulations of situations, real or fictional, in which users are able to participate. Participation and the way it is achieved are the key components which make the environment virtual, and it is generally agreed that such participation must be engendered so that users are able to feel that they are 'present' within the environment and are able to interact with the simulation if only by navigating and moving within the scene. These kinds of environment might at first sight appear very specialised although at one level, the use of any computer software embodies elements of virtuality. As software has become ever more graphic and as users have begun to interact with software using the icons of point and click, there is a sense in which all computation is taking on elements of virtual reality (VR). Predictions for the way we are likely to interact with machines in the near future suggest that VR will become the dominant form of such interaction (Negroponte, 1995).

The way VR has developed is instructive for understanding what is now possible in constructing virtual environments which embody geocomputation. The first steps towards VR involved building digital simulations as close as possible to single users, 'immersing' them within the environment in such a way that they were connected directly, through peripheral devices such as headsets, data-gloves and the like (Heim, 1997). These developed in parallel with aircraft simulators where the emphasis was on creating realistic graphic environments in which users could execute tasks as close as possible to those conducted in a real situation. Until quite recently, the emphasis in VR has been on three related issues: representing environments using 2-d and 3-d computer graphics in such a way that users can navigate within them, introducing tactile and olfactory senses into such environments as well as achieving graphic realism, and concentrating on the way such environments are used by single individuals (Rheingold, 1991). As recently as five years ago, Kalawsky (1993) noted that traditional VR was divided into the kind of immersive activity that developed to embrace various human senses, and desktop VR based on animated computer-aided design (CAD). Both activities were dominated by the direct viewing of virtual environments rather than by processes for undertaking nonvisual activities within visual environments.

However, there have been very rapid changes in the concept of virtual reality since Kalawsky (1993) charted the field. Other developments in computing, specifically scientific visualisation have embraced VR, while the movement of computation to the net is changing the idea of participation and presence within virtual environments. To make sense of all this, it is worth developing a fairly formal approach which identifies the key components of VR so that changes due to new technologies can be easily understood. Virtual environments bring

two domains together: the *digital environment* and *users* who participate in some way within such environments. Both the environment and users must be represented and of course connected together. Environment and users interact and within such interaction, various behaviours can be modelled. Moreover, the environment itself might be modelled independently of users and vice versa although to effect virtuality, these processes must come together.

A threefold paradigm of *representation*, *modelling* and *connection* is central to the construction of virtual environments. So far, most images of VR emphasise the environment, usually represented in 3-d form using state-of-the-art graphics and CAD, with users being connected to such environments through protocols - models - which enable them to sense and perhaps manipulate the environment through navigation and other forms of body movement. However using this threefold paradigm in the context of dramatically changing technologies, opens up many new possibilities for VR. For example, representation of the environment need not be restricted to 3-d; it might be any-d, from the 2-d map paradigm to its generalisation in mathematical space. Modelling is then not merely navigation or body movement but might involve any behaviour pattern one might care to think about, from the activity of 'doing science', to 'manipulating financial transactions', through to 'taking the dog for a walk' in a virtual world. Connection however is the most rapidly changing aspect of VR at present. Fully immersive VR from whence the activity emerged, is now complemented by various kinds of semi-immersion in which users are able to combine the normal material world of debate and discussion with interaction in a virtual world. VR theatres and CAVES are the most obvious portrayals.

However, the semi-immersive nature of these environments where interaction with the virtual world is embedded with interactions in the material world, is not the main focus per se. Once fully immersive VR is opened to the material world, not *single* but *multiple* users can participate. In a sense, interacting alone in any environment will always be a special case. Facilitating more than one person to interact enables a full range of social interaction to be represented and modelled. In developments of VR technology into theatres and CAVES, interaction is still direct but what is now happening as computation drifts out onto the ether is that presence and connectivity are becoming *remote*. Users can access virtual environments across networks, can be represented graphically as avatars say, and can then engage with others who are connected at different remote locations with respect to some task. These kinds of networked virtual worlds are growing very rapidly at present, providing environments in which many new kinds of representation and modelling are able to take place. The next step, already happening, is linking such remote interaction with semi-immersive and immersive VR, thus generating a veritable cornucopia of virtuality.

We will first begin to classify virtual environments with a specific focus on geocomputation. This will frame the scope of the examples that we will illustrate in this chapter, but before we select these, we will describe a variety of different representational and modelling techniques that have and can be used in constructing virtual worlds. We will then focus on four examples from the wide array of possibilities that we chart, concentrating on environments in which the three key activities of information, analysis, and design are combined into four types of geocomputation. We will begin with information-based analysis based on internet GIS with various multimedia which opens up the world of GIS to a much wider set of participants. We will then illustrate how conventional geographical (spatial) interaction models can be explored virtually indicating how optimal locations can be investigated visually by many experts. This is followed by an illustration of how group problem-solving and the design of real urban environments can be achieved remotely through networked virtual worlds which mirror the real world in question. And finally we will show how fictional virtual worlds, designed with no specific purpose in mind other than social interaction, might form rich data sets for studying how real urban environments might be explained and simulated.

2. Classifying Virtual Environments

Providing maps for any rapidly changing emergent field is essential but impossible. Virtual reality has no formal definition as we have implied. In one sense it might be seen as the emerging interface to digital computation per se but our distinction between environments and users and the threefold paradigm of representation, modelling and connection does provide a convenient starting point. It is immediately obvious that all non-digital environments have some counterpart in the digital world, no matter how realistic, and although most of these have not been developed, nor may they ever be, then this still points to a vast array of virtual environments that might someday be realised. In terms of the environments to be constructed, it is worth making a distinction between *real* and *fictional* worlds. The line between such worlds is inevitably blurred; it is hard to define any world which is simply one or the other and this merely recognises the old adage that 'truth is often stranger than fiction'. But this is important because much of the technology and progress in VR comes from fictional applications in games or in simulated situations which have no link whatsoever to any material reality.

There is another important distinction between *spatial* and *nonspatial* environments. There is a further distinction between *spatial* and *aspatial* (after Webber's (1964) place and non-place urban realms) but here all we need be concerned with is environments which can be represented in 2-d or 3-d and those that do not bear any resemblance to such Euclidean space

such as those in mathematics, those based on writing and verbal interaction, those based on sound, etc.. Most virtual environments to date, particularly those based on full immersion, are spatial, usually 3-d, where the representation is based on the material world of buildings, landscapes, interiors. Part of our quest here is to show that these kinds of world are likely to be the exception rather than the rule as VR continues to develop (Kitchen, 1998). As with reality and fiction, spatial and nonspatial blur into one another and many applications are likely to contain elements of each.

First applications in VR consisted of *single users* being immersed in digital environments, and this has been reinforced time and again in the last twenty years as the desktop paradigm has come to dominate computation in general. In fact, immersive VR has diffused quite rapidly, particularly in the games industry where the personalised headset has become de rigueur. However with the growth of semi-immersive and remote VR, the idea of *more than one user* being part of any virtual environment has developed very fast (Batty, 1997a). Indeed, it is likely that virtual environments will eventually be defined as environments which enable more than one user to participate. Of course, the social structures that these users adopt will depend upon the reasons for interacting through such environments. Any social group in the material world might interact in the virtual but here we will define four distinct types: those using VR for science, those using VR for some professional task such as design, or problem-solving, those using VR for political purposes, and those using it for leisure pursuits.

It is difficult to develop this classification of users further without referring to the purposes which VR is likely to serve. We will make the common distinction between using VR for *information* and advice, for *science* (analysis), and for *design*. All these tasks characterise some form of serious and considered directed inquiry although we must also identify leisure pursuits as having an important influence on the development of VR. Relevant to use is the nature of that use, the common distinction being between routine (often frequent) usage, and strategic (one-off) use. Most applications to date are one-off although once VR really begins to disseminate through diverse applications, this will change. VR is more likely to become routine during the next decade, just as GIS and CAD are becoming routine, due largely to the fact that the elements of representing and modelling virtual environments are quite basic functions in themselves, do not involve esoteric science, and are fundamental to interacting with any kind of computation.

The key activities in VR are *representation*, *modelling*, and *connection*. *Representation* of environments and users can be *visual* or *nonvisual* or some combination thereof. Visual representations can be of real-material systems - buildings, landscapes, interiors - or of

abstract forms such as those associated with the analysis of problems in mathematical or statistical space. 2-d map-based representations lie between the real and the abstract but as we implied earlier, the lines between these categories are always blurred. Nonvisual representations involve text, sound, numbers whose form can be visual but whose intrinsic representation is not. Mixtures of both visual and nonvisual characterise all environments and users. The most visual environment-user system involves users appearing as avatars in some 3-d space within which they can see themselves and others navigate. But systems based on pure text or even sound such as internet relay chat (IRC) comprise VR, for users are often imbued with a strong sense of presence, such environments meeting all the requirements for VR posed above.

Modelling virtual environments include several processes which simulate the way users and environments respond to each other; we define modelling as the shorthand for ways of interacting with the environment according to standard behavioural rules applicable to particular users. *Exploration* of the environment is important and this involves some mode for navigating the visual and/or nonvisual media. Being able to intelligently respond to the environment requires *analysis* which may generate other visual or nonvisual media but the capability to alter the environment in some way requires *design* capability which is central to many virtual worlds. At one level, all three of these processes are involved in the simple act of navigation, say, but in our context, we consider exploration to be prior to analysis and analysis prior to design, notwithstanding the fact that each of these is relevant in achieving any of the others.

Finally, we come to *connection*. Single users connected through headsets to their virtual worlds represent the extreme of direct *immersion* although single users without such connection technology but in the standalone desktop environment will probably always represent the largest class of passive VR users. *Semi-immersion* loosens the hardware connection where much interaction then comes from point and click in a group context, in CAVES and theatres, but the cutting edge of VR at present is coming from the net, from networked VR where users are physically in a standalone context but are virtually in the same environment through the net. This is *remote* VR but as in all the other dichotomies we have posed, it is never a question of immersion or not, for virtual environments can be constructed to draw on the best elements of each. One last point before we show how our case studies cover this wide domain: we have implied that all the VR we are talking of here is based on computers and their networks, but other digital media might be involved for VR often mixes analog and digital. Ideas of augmented reality where TV, video, and computation are embedded together exist throughout the media but here we will exclusively concentrate on the digital world, avoiding consideration of how other media might be involved.

Arraying these many dimensions against one another defines a classification within which many possible types of virtual environment exist. To illustrate this, we will show how the four exemplars or case studies that we will present fall within this classification. Let us simplify these dimensions by first examining the distinction between real and fictional applications against the three varieties of connectivity - immersive, semi-immersive and remote, and then contrasting single or multiple users of these environments against information-based, sciencebased and design-based applications. Our first exemplar (I) begins with 2-d representations of real geographical space and extends such map media to 3-d, and other pictorial forms while generalising the single user focus of GIS to many users who can simultaneously use the system. In Figure 1, we show this case study first in terms of its focus on real environments represented and modelled in semi-immersive and remote forms, and then contrast the many user orientation of such a system to its use for information and scientific analysis. The portrayal of these classes - first as real/fictional environments v immersive/semi-immersive/remote forms of connection - and then as single/multiple users arrayed against information/science/design applications- form the first two matrices which define these classes in Figure 1: we will discuss the third matrix below.

Our second exemplar (II) involves virtual urban modelling in which real geographical problems are explored in a group context in a semi-immersive environment (VR theatre) by multiple users. Their interest is in the scientific understanding of location as well as in the design of optimal locations which can only be achieved through the kind of exploratory heuristic graphics that this kind of VR enables. Our third exemplar (III) is simpler in that here we are concerned with a realistic urban world but represented as a multiuser virtual world in which remote users engage in altering that world through design. These first three examples all embrace more than one user in remote contexts based on representations of real situations. Single user VR which originated the field is increasingly less significant for once virtual environments are possible and once networks are in place, the participation of many users interacting and communicating simultaneously and in parallel becomes the norm.

Our last exemplar (IV) radically changes this context. Fictional urban environments in which many users interact remotely define virtual worlds for which a complete record of transactions and interactions exists. When this world is an urban spatial world, then it becomes possible to use this data to simulate that world - to stand back from the virtual environment and to attempt to understand it in conventional terms - or even in virtual terms. Thus we move the entire modelling process to a meta level. No longer are we concerned with modelling a real or fictional world virtually, but with modelling the process through which the virtual world is constructed/modelled/simulated. There is an element of infinite regress in this recursion whereby we model digitally what is modelled digitally Our final exemplar thus moves

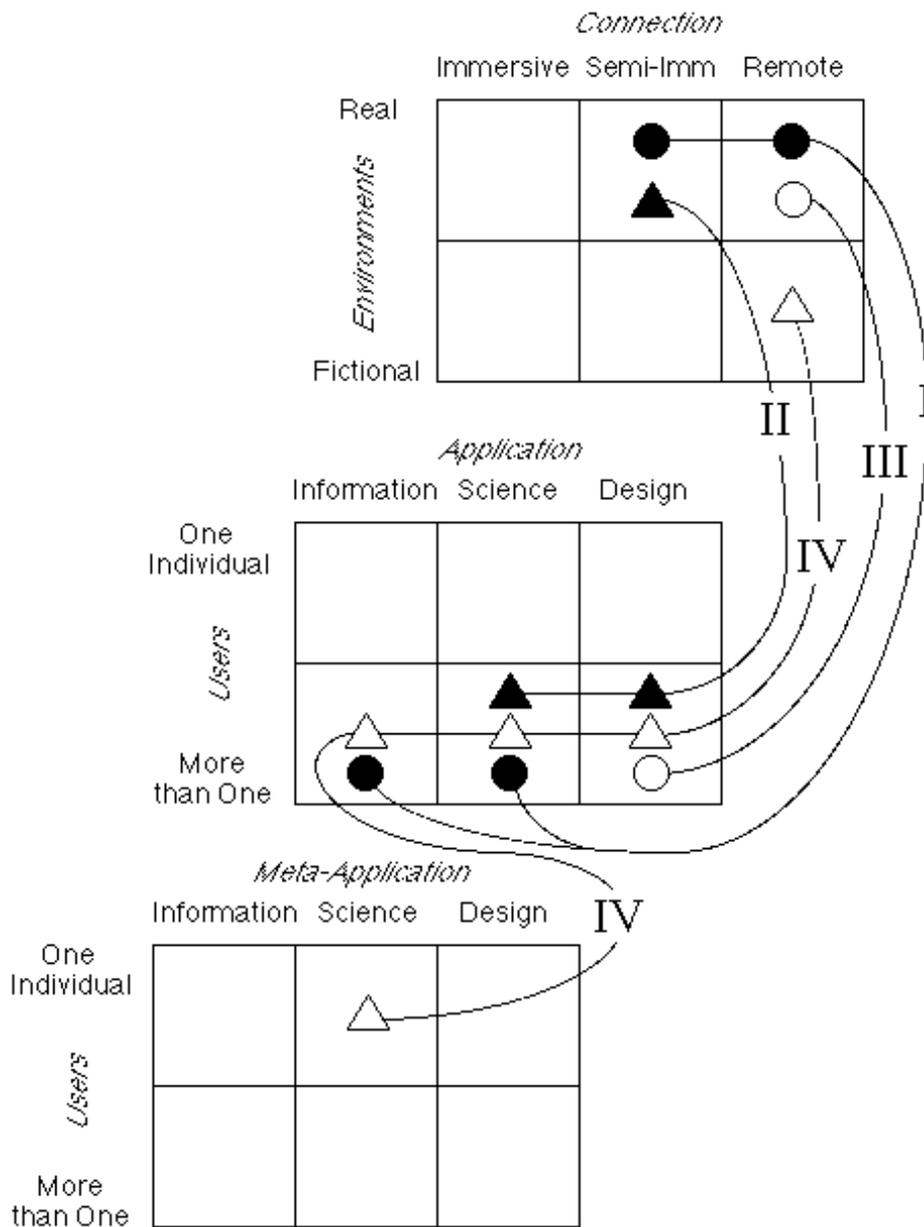


Figure 1: A Classification of Virtual Environments

to this meta level as the third matrix in Figure 1 makes clear. But before we illustrate these, we need to explore the ways in which virtual environments might be represented and modelled.

3. Representing and Modelling Virtual Environments

The majority of virtual environments to date have been represented graphically as 3-d scenes, often rendered in considerable detail so that users are given an immediate sense of visual realism. Users of such environments usually appear in visual terms as avatars, the term used to define a virtual representation of an animate object (from the Hindu definition of the 'descent to earth of a deity in human, animal or superhuman form' - OED, 1991). Avatars which are usually scaled to the 3-d scene are more commonly used to represent many users in remote contexts whereas in immersive VR, the emphasis is on the representation of single users in human form which is closely matched to eye, hand and other body parts.

Modelling in such contexts merges with representation. The representation of 3-d scenes is often referred to as modelling with the scenes themselves as 'models'. For example, large scale block representations of cities are called urban or city models, not to be confused with the more generic use of the term modelling as mathematical and statistical simulation which characterises mainstream urban and geocomputational modelling. When users are introduced into such scenes, the emphasis has been on enabling them to navigate through and to manipulate objects within such scenes. These kinds of processes where real behaviour is mapped onto the virtual context is clearly 'behavioural modelling' in its widest sense, building on insights and theories from graphics, cognition, way-finding, and related areas of psychology and human-computer interaction.

However in this context, we are much less interested in these kinds of urban environment. As we have already noted, VR is rapidly moving into network and other visual and nonvisual contexts where much more purposive action and decision with respect to use is being developed. Environments and users based on 3-d representations are likely to be the exception rather than the rule; VR is broadening to embrace users who do not appear formally in those environments but whose actions are outside the environment *per se* and whose behaviour patterns are much more subtle with respect to their interaction with the digital environment. A few virtual environments exist which are 2-d map-based but their structure is not dissimilar to the norm of the 3-d environment with the exception that users usually remain passive: their location in 2-space may be visualised but the kind of equivalence between users and environment - avatars in 3-d scenes with full motion - is not possible. In short, the traditional paradigm of VR based almost exclusively on users within scenes has little meaning once the move into other visual and nonvisual environments takes place and as other forms of information and connection become important.

What is required is a broader and richer paradigm which enables the various elements of virtual environments and their use to be discussed. We do not intend to develop a fully-fledged paradigm here but we need a framework for considering the array of different representational and modelling styles that might characterise an expanded definition of VR. To this end, we consider the idea of linked information to be an important way of approaching VR. If we define VR to be the kind of environment in which users 'feel' a sense of being part of the problem with which they are concerned, in which they are able to develop a 'deeper' and more thorough understanding of such a problem through analysis and manipulation, then this casts the VR net much wider than the simple notion of working with visual environments that resemble scenes where the focus is simply upon navigation. Our augmented definition would thus include virtual laboratories, virtual lecture theatres, web forums, and a whole host of digital simulations which would not be visual in terms of their material origins. The interface to these environments would still be largely visual but the processes of exploration and manipulation might be distinctly non-visual. The virtuality in this sense would come from representing the environment through many different kinds of information simultaneously available to the users, so linked that users are able to respond to an environment rich enough for them to gain real insights into the problems in hand.

The idea of linked information covers both visual and nonvisual information. Visual information embraces multimedia in such a way that environments might be characterised by a subtle mix of live and canned photo and video, structured textual information, maps, 3-d models of various kinds from animated scenes to fully-fledged CAD models. Nonvisual information embraces other senses but in particular, textual and numerical representations and manipulations are central to such VR. The ways in which these representations might impart a sense of virtuality is likely to be through various ingenious collages and juxtapositions of the information but the concept of 'hypertext linking' or 'hotlinking' different information is central to the way these virtual environments might be constructed. In such environments, users could be actively represented - as avatars in 3-d scenes or their presence might be portrayed in 2-d. But in general, their representation is likely to be passive, to be configured as part of the sequences of using such environments which may be recorded in some way but not necessarily visually represented within such environments.

However, the real power to interact in such environments would come from a very wide array of modelling types which might be activated by users. Four kinds of process suggest themselves: *navigation protocols* for traversing and exploring such environments in the broadest sense; *decision protocols* for interacting with other users and reaching agreement or otherwise over common problems and goals; *analytical functions* for manipulating information, using a variety of formal and informal 'scientific' procedures, many of which

might be part of conventional desktop software; and *manipulation functions* which deal with ways in which information might be changed and combined and new information introduced into the environment: design, for example, would form one type of such function. We will deal with these in turn.

Navigation protocols do not deal simply with movement in 3-d space but with the entire way in which users explore the information contained within the virtual environment. For example, a user may literally move across a map, accessing 3-d scenes, photos, video clips, textual information, sounds and so on, and it is the overall process of access - the way that access is initiated, the design of that access, and the way in which users juxtapose information and generate insights - that is part of this domain: in short, this is the human-computer interface but structured with particular purposes and problems in mind. There is a clear design task here in that in modelling such environments, the best ways in which such navigation leads to insights, must be constructed, and this will involve much trial and errors with users, and with the representation of the environment. When many users are navigating an environment, then these protocols must extend to ways in which user paths cross, and this leads us directly to modelling protocols based on interaction and decision. Note as well that part of the process of developing such environments involves logging and charting navigation so that the history of the ways in which users have interacted with their environment can be used positively to direct further interaction.

Decision protocols involve ways in which users interact to some purpose. Like navigation, interaction can take place on different levels. Many users might be present in a virtual environment but their interaction may be passive. They may not even be aware of others as in the case, for example, of several users being logged into the same web page. If users are aware of each other, their interaction may be entirely casual. The 'ActiveWorlds' server software illustrated below contains many virtual worlds whose sole purpose is to engage in virtual chat, where the designers have not provided any functionality above the level of 3-d and avatar representation other than the ability to chat to other users who appear in the scene and who are able to articulate basic gestures such as frowning, waving etc.. These kinds of interaction can be developed in non-graphic ways as in early versions of multi-user domains (MUDS) and IRC but what we have in mind here is much more purposive.

The interaction that we are concerned with involves the analysis and design of urban systems and thus interaction that is important must involve extracting various abstractions from the urban scene. A good deal of this might be concerned with single user interaction with the environment itself but when users are required to jointly engage in some task, then the interactive environment must be closely specified. There is very little research in this area as

yet. Existing psychological and sociological research on non-virtual problem-solving environments is difficult to transfer to virtual environments. There is some work on collaborative spatial decision-making using GIS which heralds this kind of interaction (see Carver, 1997) but there are few case studies as yet. Although many virtual environments based on multiple users are being developed at present, the emphasis is on use and application, rather than on understanding the best ways of reaching decisions. Clearly charting decisions, developing ordered structures in which information can be abstracted and interpreted, facilitating internet conversation and so on are all issues that have to be reflected in the design of such environments.

Developing basic functions which users employ to generate insights is more straightforward as many of these already exist as part of software which is being used to construct virtual environments. Traditionally VR systems have hardly any *analytical functions* which users can employ to make sense of their environment because the emphasis, as we have seen, has been upon 3-d representation and modelling through visual navigation. However in an analytic context, then the complete range of functionality that exists in desktop and workstation software can be adapted. Already 2-d geographic information can be distributed across the net and environments constructed which access this (Plewe, 1997). GIS software is being adapted to multi-users, and web forum software enables users to begin to connect up to such data in a structured way. In our first case study on virtual GIS, we will show how simple desktop functionality relevant to abstracting and modelling 2-d map data is already available in a remote context, while adding analytic functionality to CAD in a web context is under way. In fact, the development of 3-d analytic functionality based on web software such as VRML, the development of platform-independent applet software and such like is well advanced. During the next 5 years, there is bound to be an explosion of applications stretching the concept of virtuality along the lines we have speculated here.

Our final set of modelling tools involve *manipulation functions*. Manipulation in traditional VR involves reconstruction of the 3-d scene in various ways by users from within the scene but the kinds of manipulation we have in mind here are much broader. For example, if the environment were based on some mathematical model, then manipulation might involve sensitivity testing, or control/optimisation of the model. There might be some visualisation of the solution or phase space which is tied to manipulation as we will show in the virtual urban model, the second of our exemplars illustrated below. However, there are many tricks which we can use to visualise such manipulations but the key issue is that considerable functionality is likely to reside behind any visualisation. In traditional VR, this has rarely been the case but for example, even in manipulating the components of 3-d scenes, there is a need for evaluation functions which display how the manipulation meets prior goals. In turn, these

must be specified and appropriately represented within the virtual environment. In short, manipulation of any kind must be to some purpose and therefore is likely to involve the use of functions which indicate how this purpose is being achieved. Finally, all four of these modelling functions - navigation, decision, analysis and manipulation - cannot be easily separated from one another for to implement one is only possible through the others. In the four exemplars that we now discuss, we will indicate the essential nature of this interdependence.

4. Exemplar I: From Internet GIS to Virtual GIS

We are constructing an internet GIS for Greater London which is being linked to a variety of other visual and non-visual information. In Figure 1, we classify this environment as existing in semi-immersive and remote form, accessible to many users for purposes of information and scientific analysis. The environment can be accessed in a VR theatre on the full screen with many users able to interact with the software in a group decision-making context but with other users logged onto the software remotely. At present, the heart of the system is the distributed GIS from ESRI called ArcView Internet Map Server (AIMS). This software exists on a web server and whenever a client (remote user) delivers a request to the server, the server fires up ArcView which processes the request and delivers the output back to the client as a Java applet called MapCafe. MapCafe has the look and feel of the desktop GIS called ArcView; AIMS can be tailored to deliver considerable functionality characteristic of conventional GIS to remote users. For example, layers can be switched on and off, zoom and pan features are standard, query functions are central to the software while several other functions involving 3-d visualisation, network analysis and spatial analysis can be invoked.

Although the GIS is at the heart of the environment, because it is accessed through a web browser, then all the hyperlinking associated with the web is possible. We have exploited this through hotlinking the GIS to other visualisations of the data, namely video clips, related text and numerical information in basic tabular/page form, VRML models, and other web pages. Because the model is in its earliest stages of construction, we have not yet begun to develop ways in which users might be represented within the environment, nor have we examined how any formal interaction between them might be structured and recorded and protocols devised for reaching decisions. The ultimate idea of the system is that information will be available for several different types of user who will have different sets of privileges enabling them to perform analyses, store and represent their own results within the system, browse the system, converse with other users for specific purposes, download data for their own use, upload data into the system, and move towards the design of planning and environmental

policies at different scales. The system will ultimately link scientists, policy makers and specific interest groups as well as being available for more general public access.

The system only exists as a demonstration within the VR theatre context as yet but to give some idea of its potential, we will illustrate its current working. The server which runs the internet GIS is mounted on a highly specified PC which also contains the web browser software through which it is accessed. The VR theatre where the system is viewable in its complete form is driven by a Silicon Graphics Reality Engine (Onyx 2) with two channel output which generates a double size screen with dimensions in the ratio of 12:5, the scale of the entire end wall of the theatre. The Netscape viewer is configured accordingly. The interface is in two sections. The output of the GIS as MapCafe is presented on the left hand side of the screen, always in 2-d map form; on the right hand side of the screen, various information is loaded relating first to the map and accessible as hotlinks from MapCafe - VRML models, videoclips, photos, text pages and other web pages. However, from these outputs which are accessible only from the map, further information is then displayed in the same area of the screen from hotlinks associated with this other non-map information. So for example, the user can hotlink from the map to a VRML model but then from the VRML model the user can hotlink to videoclips, other web site and so on. The map always remains on the left of the screen but the information on the right may eventually be many levels removed from the map as the user searches other information within the system as well as the web. The navigation is thus fairly basic with the map always as the starting point and with the user always moving back to the map to initiate any new search through the trees of information accessible within the browser.

To give an idea of what is possible, we show the scale of the browser in Figure 2. The left window is MapCafe which resembles ArcView but is in effect the Java applet which AIMS delivers in response to a query from the client. The layers in MapCafe (seen in legend on the left of the map) can be switched on and off in usual desktop GIS manner and as the user zooms, different layers of information can be made active. In the scene shown which is Parliament Square, then the base map is Ordnance Survey Landline data which is laid over the Cities Revealed ortho-photographic data base for London. In that scene, it is possible to load a VRML model for the entire area of Central London by clicking a hotlink in the corner of the Square. This model is loaded into the Cosmo Player which is shown on the right of the screen. It can be rotated and viewed from any angle; that shown is looking into Trafalgar Square with the Thames and South Bank in the distance. This model has been crudely developed using the 3D Analyst extension to ArcView and as yet has not been rendered to the level of realism that might be expected. It is possible for the user to enable AIMS to generate

any VRML view from the OS Landline data from within the internet GIS because this is now standard functionality in ArcView, subject of course to appropriate licensing.

When a user first logs on, the extent of the London data is first displayed. Figure 3 shows the left hand screen map which appears on entry. This is a map of the London boroughs with different sizes of town centre shown as proportional circles. In the demo, we have constructed so far, it is our intention simply to illustrate possibilities and thus we have incorporated diverse data of this kind. Also shown on this screen in the centre of London are two boxes; the first defines the area where the user can zoom into the Bloomsbury district where University College London is situated and to explore a VRML version of the Quad, to load related video scenes and to access various information about the College and the area; the second enables a zoom into central London (Westminster) and to access the scenes shown in Figure 2.

To illustrate the functionality and diversity of the system, we have layers of information in the system that pertain to health care from the Riverside Health Trust which covers the Hammersmith, Kensington, and Ealing areas of London. When we zoom into these areas, MapCafe presents medium scale road and rail network data (Bartholomews) and onto this we can switch relevant layers. A contour surface whose density is proportional to the number of visits by health care workers to patients is shown in the zoom in Figure 3, and the spider lines show how we can reallocate health care workers from their base stations to patients in such a way that travel cost is minimised. This has been done within ArcView using the standard Network Analyst software. It is illustrative only of the kind of data that can be preprocessed and displayed, but this implies the range of functions that ultimately we will enable for many remote users within the internet software.

So far, the navigation protocols for this system are elementary. We have not yet embedded the product into a web forum and thus the decision protocols are barely developed. However, use of the product in the VR Theatre context has been extremely successful, simply in fashioning the demo and in engaging debate about the whole idea. The functions so far are all those which are possible within internet GIS and as yet there is no capability for users to produce policies or designs. However, in the various exemplars which follow, the kinds of design capability that we will illustrate there, can and may be used as this virtual GIS is further developed.

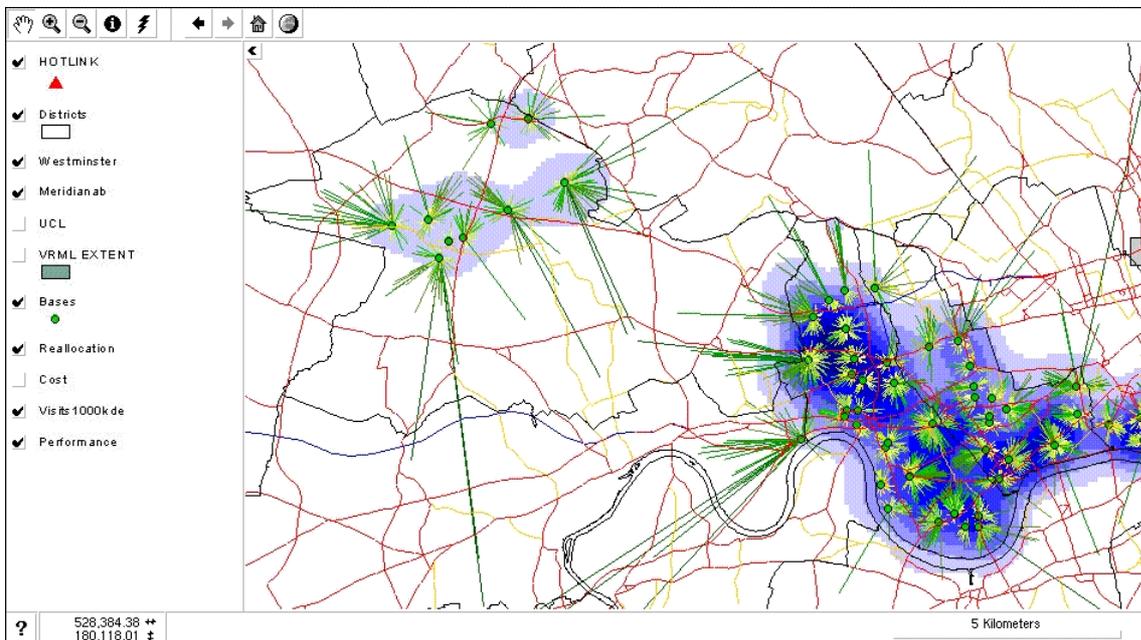
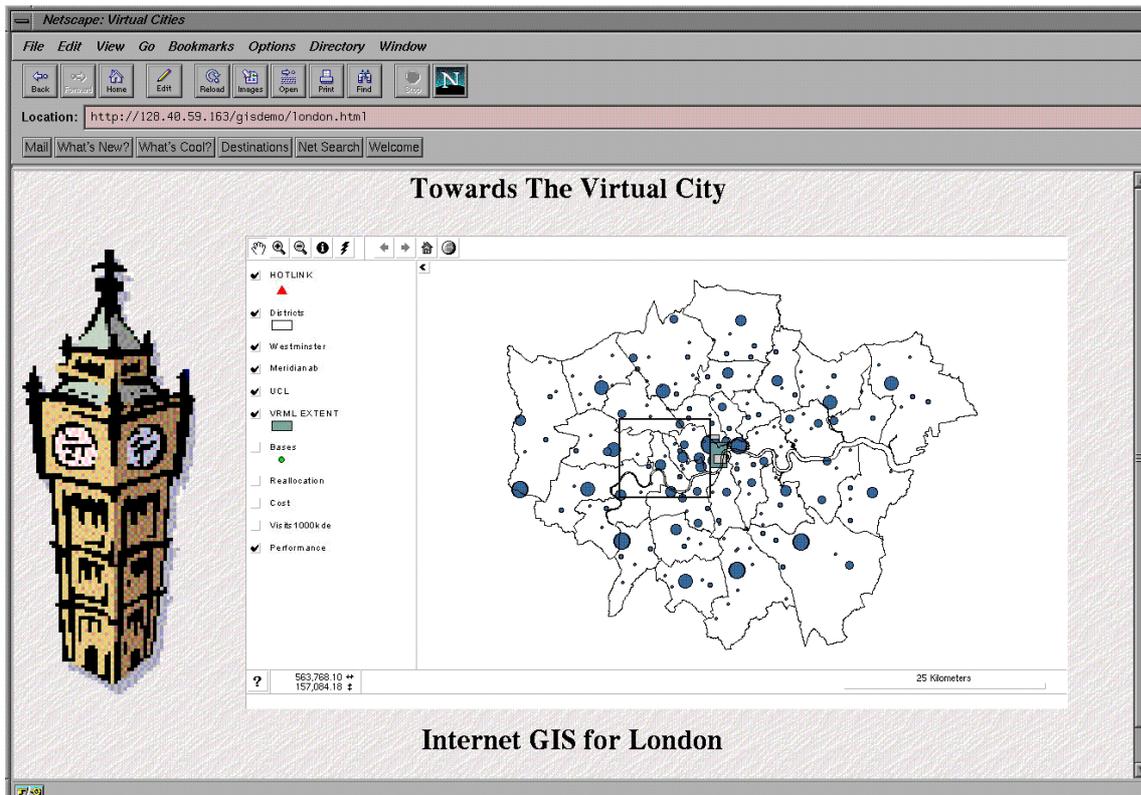


Figure 3: Zooming into and Displaying Different Layers Using MapCafe

There is a long tradition of spatial interaction modelling which is one of the cornerstones of geocomputation. Gravitational models have been used since the 1950s for all kinds of locational and transportation planning and considerable effort has been put into an understanding of their mathematical structure. As part of this, these models can be viewed as processes in which individuals optimise their location. General frameworks have been developed which show that the resultant patterns observed in cities are consistent with accessibility-, entropy- or utility maximising where the objective functions of access, entropy or utility are conceived as spatial averages describing the dispersion of tastes within any population. This dispersion effectively means that everyone does not locate in the same place and that the kind of distance decay patterns that characterise the way individuals interact and travel in real life are effectively replicated (Wilson, Coelho, Macgill, and Williams, 1981).

The fact that such models can be easily placed within a framework of optimisation means that there is the chance that optimal solutions under different constraints might be found for a variety of practical location problems in which equity and efficiency considerations can be traded off against one another. However, these kinds of model are intrinsically nonlinear; they can be solved exactly in their pure form where there is a global objective function of the kind used in their derivation but once more realistic and practical constraints are introduced, such models become impossible to optimise directly. For example, consider the problem of a retailer or developer wishing to locate a shopping centre so that his/her profits are maximised. This problem has been handled in the past by developing a spatial interaction model which links potential shoppers at their home base to shopping centres, and exploring how much trade can be captured when new shopping centre locations are introduced into the retail landscape. This problem is usually solved by trial and error: plugging in a shopping centre and evaluating how profitable it is relative to other locations. Because there are often only a limited number of locations for new centres, then each can be evaluated and their profitability assessed.

There are versions of this problem that can be set up as formal optimisation but in general, even with relatively simple problems such as these, it is not possible to solve the optimisation problem directly by using the model to mathematically derive the most profitable location for a centre of a given size. Real problems of course are much more complex. Retailers and developers face an environment in which the decision is not just to find the best location for a new store. They might expand, contract or close existing stores; they might develop different sizes of store for different market niches; they are faced with other competitors attempting the same, with a variety of planning constraints which affect what they might do, all against the vicissitudes of the market place and capital finance. They are also faced with the fact that what they might do in the short term, influences the medium term. Locating a new store

might attract others to it and this kind of externality affects what is profitable over different terms. In short, such problems are highly nonlinear, full of positive feedbacks in space and time which make their solution mathematically intractable. The only way forward is through the use of heuristics to get some sense of the solution space in which such a problem exists. Virtual reality systems offer much promise as we will now illustrate.

Imagine that the quest is for some developer to maximise retail turnover by locating one or several stores within an existing retail landscape. The model which drives this problem is not one with simply a spatial dimension but it has a temporal dimension in that once a store is increased in size or a new store located, then there are positive feedbacks in that the store attracts others to its location or it begins to fail. The time taken to develop the store is also critical, so the problem is one that must be solved both spatially - in terms of location - and temporally - in terms of the phasing of development. This problem resembles those formulated by Clarke and Wilson (1983). There is an immediate visualisation of the geographical retail turnover surface which will clearly change as centres change in size and location. However when a developer examines this surface, then the location where there is maximum turnover is no longer necessarily maximum once a centre of a given size is located there. There are competition effects as well as positive feedback effects which change the surface. Add to that the kinds of constraints that characterise real conditions such as the fact that locations may be places of maximum turnover but with no space to develop or the land is owned by those who will not sell or the cost of land is too high, then these must also be built in as constraints on the optimisation. These are often hard to describe in the kind of mathematics that optimisation requires for once highly nonlinear problems are subject to very mixed constraints - integer, binary and so on - then their chance of solution is almost non-existent.

Visual solutions or rather visual solution sequences in which users can move towards better solutions are the only way forward. We are developing such an approach in a semi-immersive context. In Figure 4, we show a map of the town of Wolverhampton where the central area (within the ring road) has been set up as a 3-d VRML model (generated from ArcView). Draped over the entire map is the existing retail turnover surface (also generated using Spatial Analyst within ArcView) from data on retail turnover at unit postcode level (~100m resolution), made available for a related project by the Office for National Statistics (DETR, 1998). This visualisation uses Cosmo Player to view a VRML model within a web browser and when displayed in the semi-immersive context of the VR Theatre, it becomes immediately possible for groups of interested experts to examine the meaning of the surface and changes to it. The following kind of visualisation and optimisation thus becomes possible. The turnover surface is generated by a highly nonlinear retail model calibrated to

the existing situation. The inputs to this model are size and location of centre, travel network and the density of consumers (population), all of which can be visualised in the manner of Figure 4.

As these inputs are changed, the surface will change and thus the environment we have under construction is one in which the inputs to this model can be manipulated directly by users who can then view the outputs of the model directly on the screen. Size of centre can be indicated in the 3-d block model, density of population we can assume is fixed, but the route network can also be manipulated graphically. Assuming that only the size and location of centre or store were to change (the developer's location problem), then it is possible for the user to drag a centre to different positions in the map and as this process takes place, the model is computed on-the-fly and the surface redrawn. At SG Reality Engine refresh rates (computation at 10 million polygons per second), the surface changes at speeds in excess of the time it takes to move the mouse and thus real time motion is engendered. The user can also change the size of the centre using various sliders and the program is so configured to stop the user infringing any basic constraints where location is forbidden or not feasible.

Optimisation is achieved visually. As the location and size of the centre changes, so does the retail turnover surface and what might look like a good location at one instant may not be once the centre is dragged to it. To make progress in moving to a better location, then a variety of other information must be displayed and computed. For example, the derivative surfaces can be displayed, showing best directions of search and the entire armoury in visualisation of phase spaces employed to illustrate progress. Of course, this visualisation has assumed that the objective function is retail turnover per se but there may be other related measures which cannot be visualised geographically. In fact, this problem needs to be researched as a problem in scientific visualisation as well as one in virtual optimisation, and to this end, it likely that in the ultimate interface will not only consist of a geographical phase space as in Figure 4 but more abstract windows on the model as well as sliders and other devices for controlling the optimisation. We are exploring a variety of VR software appropriate for this implementation.

Finally we should note some features of the modelling styles used here. There is extensive functionality that users would make use of here including many GIS-like functions based on statistical and modelling procedures; these would be linked to ways of manipulating the model visually, mathematically, and through its optimisation. The 2-d and 3-d representation is fairly standard and there are well developed techniques as we show in Figure 4 for these but the decision and navigation protocols are not really specified. This environment so far relies on the expertise of users seated in the material world - in the VR Theatre or in whatever

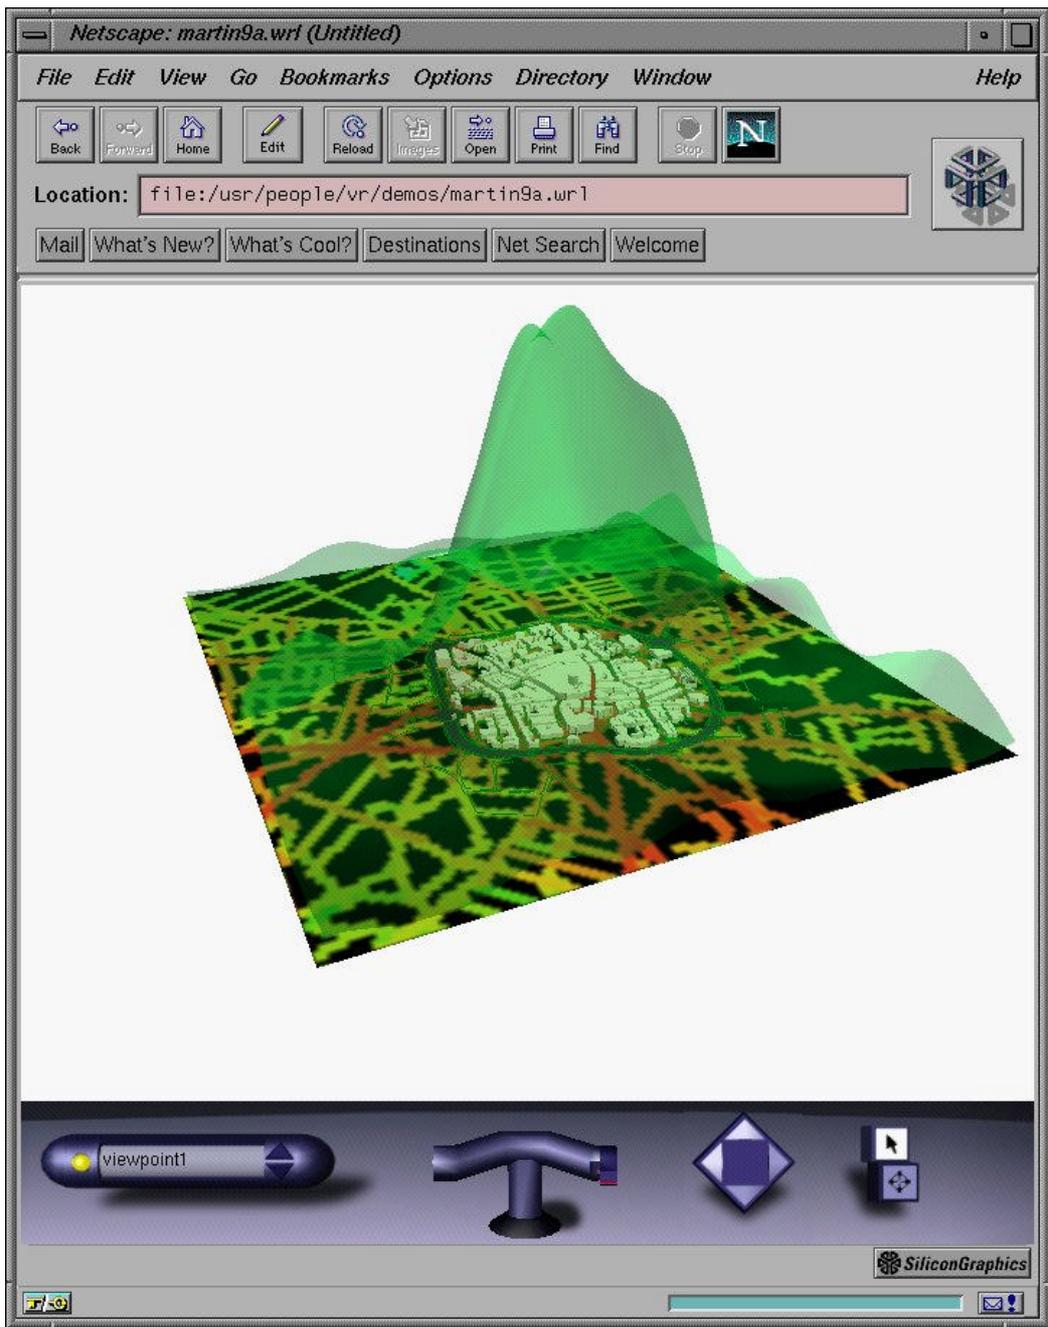


Figure 4: 2-d and 3-d Visualisation of the Retail Location Problem

environment this type of software is displayed - although there is clearly considerable work to be done of the way users begin to respond to the optimisation process. This is another illustration that virtual environments are so new that until we begin to use them, we will have little idea as to how to design them in the most effective way.

6. Exemplar III: Virtual Urban Design

The most well developed virtual environments in the visual sense are those in which both environment and users are visually configured with users remote from one another materially, but together virtually. These are those virtual worlds which are sprouting up everywhere on the net (Dodge and Smith, 1998). We will illustrate what is possible from AlphaWorld, one of the largest worlds (circa 200 000 users) on the ActiveWorlds server, but before we do this, we should set the context for urban design. One of us (Andy Smith) has developed a 2-d to 3-d virtual urban information system which enables users to tour selected sites in central London by navigating from points on a map through 3-d scenes captured as wrap-around panoramas from which access to other relevant web sites can be made. The software is called Wired Whitehall and one version hangs off our webpage at <http://www.casa.ucl.ac.uk/> (see Batty, 1997a). The idea of hotlinking map sites to 3-d scenes is a basic tool which we are using in our extension of internet GIS to virtual GIS which we presented above, but Wired Whitehall is now being further developed using techniques and ideas which are used in virtual worlds software.

There are three basic components of a virtual world: a 3-d environment which is represented in VRML-like form but which can also be manipulated by users; representation of users who log on remotely as avatars and who are able to interact with each other through primitive visual gestures; and windows in which users (as avatars) can converse through text appearing as script which is a record of all that passes ‘conversationally’ within the world. Such a world is illustrated in Figure 5. The chat box and the window which list those within the world are clearly visible below the 3-d representation. Three avatars are shown in the scene and the record of conversation is typical of casual interaction within such worlds. But the ability to manipulate the 3-d scene is also shown. A user can only manipulate or design on a plot which he or she owns and this involves being a fully paid up member of such a world. Basic manipulation of 3-d shapes is very easy: on AlphaWorld, Andy Smith’s Online Planning Office was and is under continual construction, except during the period when his subscription ran out and the plot was immediately ‘destroyed’ by the owners of the world !

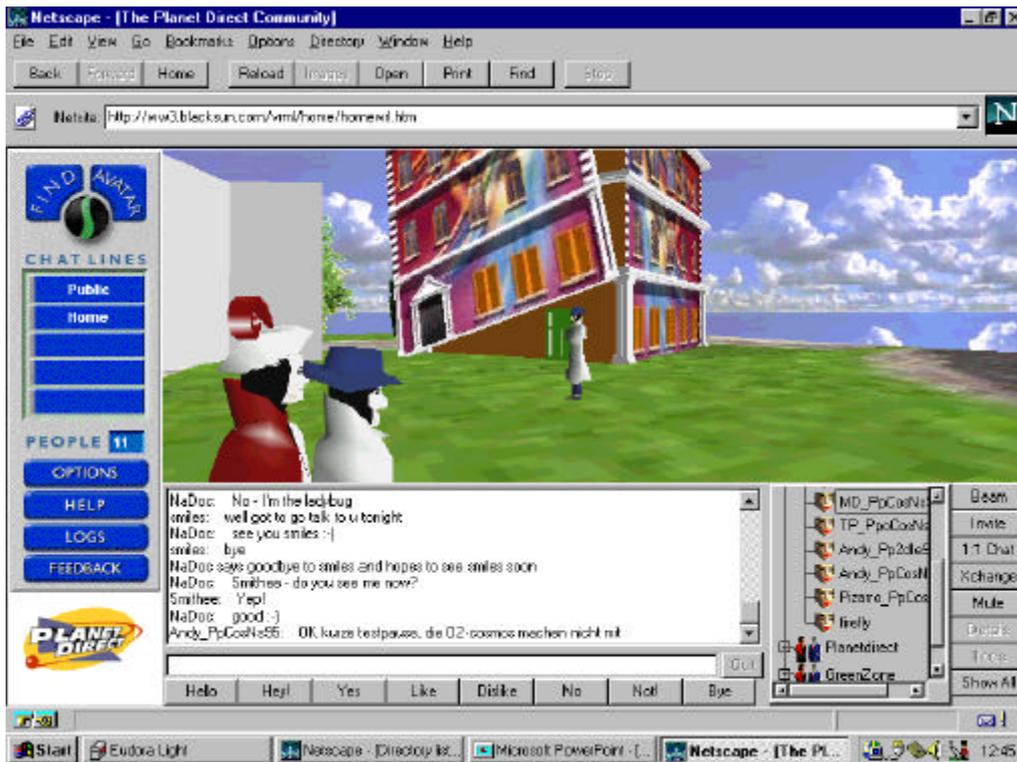


Figure 5: Avatars Manipulating 3-d Built Form in a Virtual World

Meshing this kind of environment with the kind of interface in Wired Whitehall, involves making the 3-d representation in the virtual world considerably more realistic. It is possible to combine digital photo and video with objects which can be moved within scenes. A good example of this is shown in Figure 6 where a panorama of residential development around the Surrey Docks, activated using the same kinds of software used to enable its animation on the web (as in Wired Whitehall), is displayed together with a ‘bouncing ball’. The ball which bounces all the time can be positioned anywhere within the scene by the user and this contains all the elements necessary for certain types of detailed urban design. Andy Smith is developing the Wired Whitehall virtual tourist into a virtual world in which any user who logs on and navigates into the same scene as any other user, will appear as an avatar, can engage in ‘chat’, and can move certain objects within the scene, even if these objects are not under their individual ownership. This is the basis of virtual urban design. If the chat is purposive, to locate features within the urban scene which represent the best designs that those in the scene can develop, then design can take place virtually. An illustration of what we are moving towards is shown in Figure 7 where the object of interest is a telephone box to be located somewhere within the 360 panorama at the junction of Whitehall with Whitehall Place. A combination of cursor keys, mouse and menu items enables the user to control his or her own motion in the scene as well as the motion and location of the telephone box.

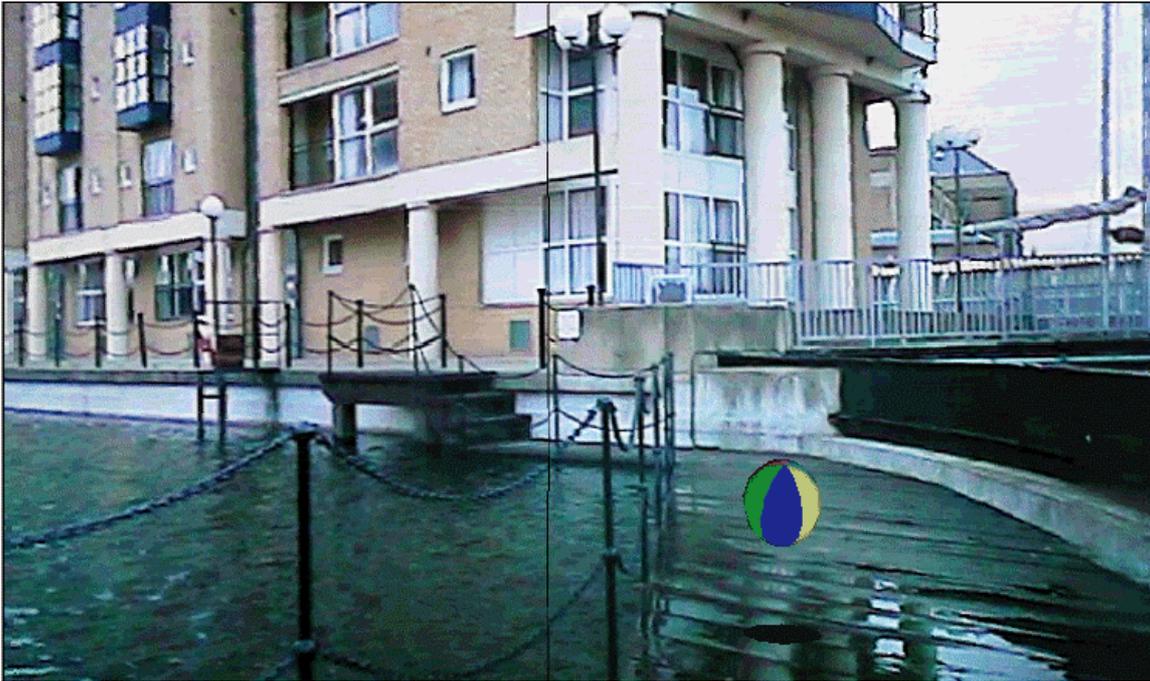


Figure 6: Manipulating a 'Bouncing Ball' in a Photorealistic Panorama

There are several features of this interface that require comment. The navigation protocol is straightforward and is that adopted as standard in virtual worlds software. This is not something that we intend to change very much but the placement of objects within the scene to optimise design is a feature that we are developing. However because these manipulation functions involve design, we are considering how we might develop some more abstract characterisation of what constitutes good design in such a world. At present, the evaluation is entirely subjective, based on agreement. Moreover, there is no necessary basis for agreement other than the fact that presumably two or more users must agree that there is a design problem to be solved. We may decide to incorporate some stronger rules into the interface to achieve some convergence but all this implies that there should be different pathways to decision that users can invoke, one of which must always be 'no' agreement. As yet, there are no analytic functions which users can generate in their quest to design but we are considering adding features which analyse scenes from a qualitative viewpoint in terms of perspective views and such like. Text-based information can, of course, be used to supplement such visual interfaces. But we see the greatest progress being made by synthesising real video and photo with VRML models through which avatars can move and from which we can extract and display to users basic information about the environment's geometry.



Figure 7: Virtual Urban Design in Wired Whitehall

7. Exemplar IV: Modelling the Morphology of a Virtual World

Our last example is very different. The three exemplars so far all illustrate how users can begin to generate insights in geocomputational terms about the urban world by entering those worlds, engaging in dialogue, extracting information, and manipulating the worlds through dialogue and intelligence derived from information. All three of our virtual environments have these features to a greater or lesser degree. However there are a set of problems of understanding virtual environments that exist at a meta level. For example, when we inquire into the best ways of designing protocols and functions, then these types of human-computer interface problem are at one level removed. Yet most of these meta problems are not urban ones per se although there is one meta problem which spins off from these kinds of environment which does involve urban modelling. Many virtual worlds are urban worlds in cyberspace, but sufficiently like the real urban world to be subjects of understanding in their own right. Our three exemplars do not fall into this class; they are already abstractions of the real world into the digital and there is little purpose in thinking of these as urban worlds to be studied in their own right for much of what characterises these environments is designed to

elicit insights into the real world. But urban worlds which are entirely fictional are worthy of study (Macmillan, 1996). Our last exemplar introduces these.

Alphaworld, as we have noted, is the biggest virtual world we know of, certainly the biggest running on the ActiveWorlds server. It is essentially an urban world where the emphasis is on the ownership of land and where a city structure has clearly grown up. There is no real distance in AlphaWorld in that it is possible to teleport anywhere within this world, and thus one of the most distinctive structuring factors of real urban space is missing. However AlphaWorld appears to have grown as real cities have grown. Its morphology as computed in December 1996, is shown in Figure 8 where it is clear that it has a distinct CBD (ground zero), and distinct axes in the 8 compass quadrants where development has clustered. From the Active- Worlds homepage (<http://www.activeworlds.com/events/satellite.html>), the owners of the world say:

“You can see the ‘Starfish’ shape of the building as people crowd their buildings along the north-south axis and ‘equator’ of AW, and as they build along the coordinates with matching numbers. (ie 220n, 200w, 450s, 450e etc.) Some do this so that their coordinates are easy to remember, and others simply by building onto what others have built”.

This is a fascinating quote in that clearly these factors - imageability and agglomeration economies - are important but perhaps AlphaWorld reflects the quest for space too - for isolation as well as congestion, the two forces which define the tension of the modern city. The challenge is thus: can morphological models based on simple processes of attraction and repulsion, concentration and dispersion, be built which reflect the way AlphaWorld has developed ? And can the processes of decision which define this growth and change be linked to such models, thus establishing the relative importance of the various urban forces and their relationship to different populations types ? Are there several processes at work and is there spatial differentiation of these within the morphology which has developed ?

What is remarkable about these virtual worlds is that the morphology mirrors in some sense real cities, but unlike real cities, there is a complete record of what has taken place in these worlds from the server log, if it is not destroyed. This provides a potentially rich source for not only mapping and understanding urban cyberspace but also for beginning to speculate on whether the kinds of processes that take place in fictional worlds throw insight on processes taking place in the real world. In modelling real urban morphologies, distance is all important and the entire heritage of urban theory from central place theory and spatial interaction to fractal geometry is built around the notion that location can be explained as a trade-off

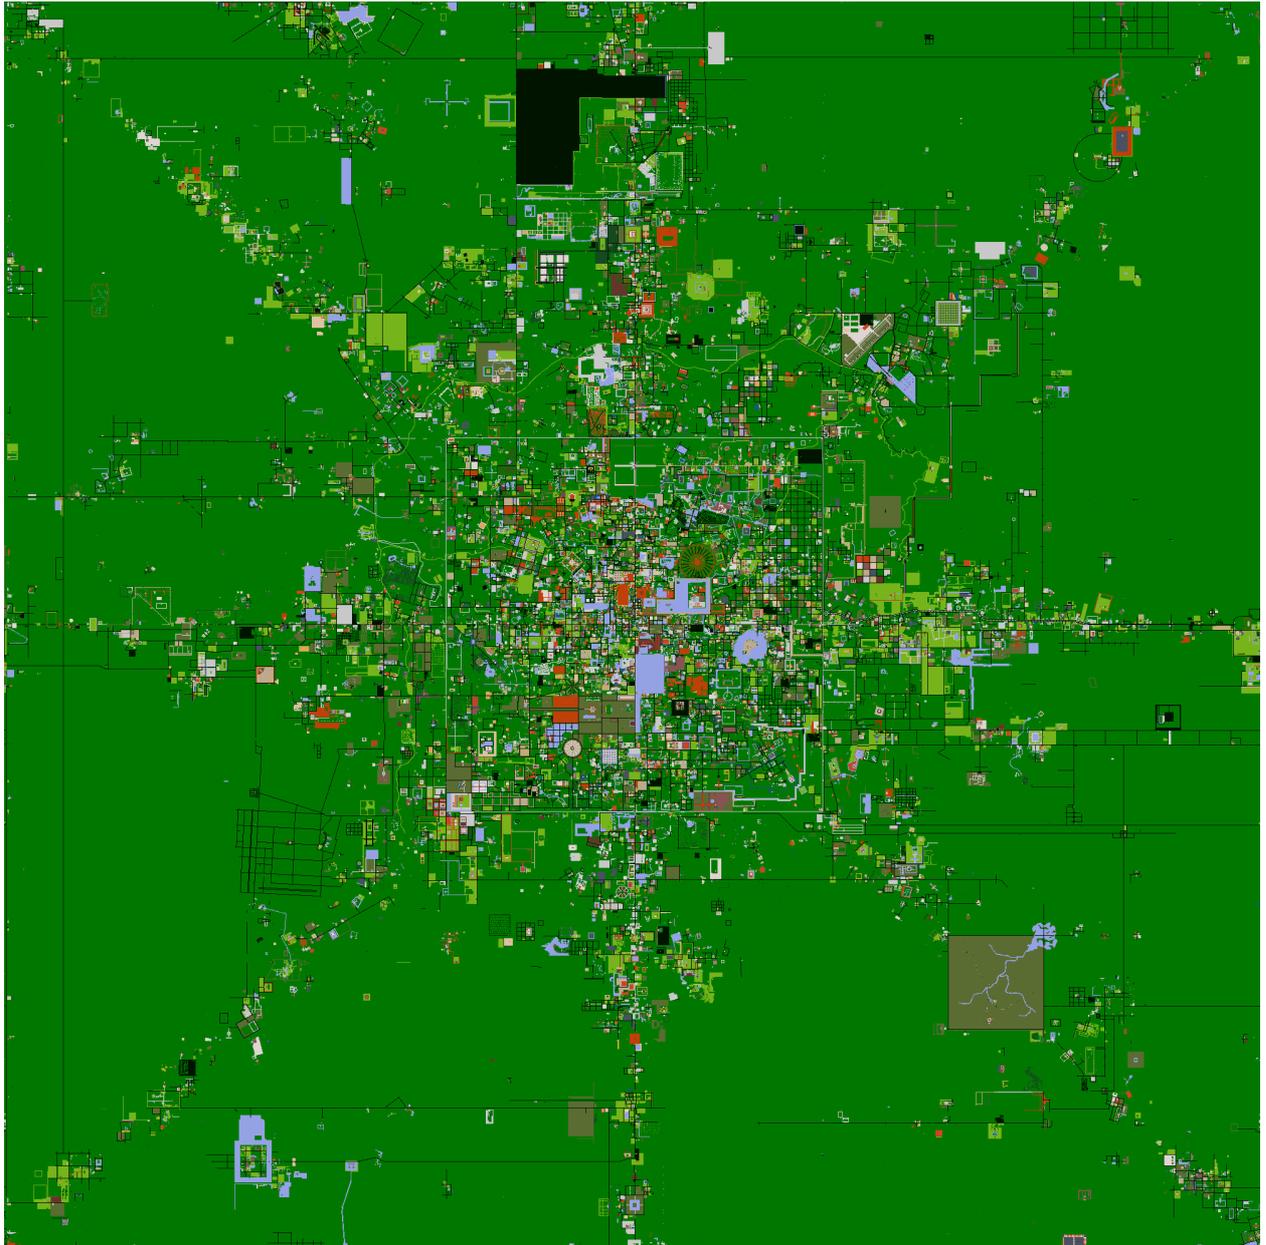


Figure 8: The Morphology of AlphaWorld

between mass (size) and distance (travel cost). In operational models that have been built, one of the features that has been examined and has always proved curious is that if the distance effect is relaxed entirely, thus no longer relevant, then these models still generate morphologies that are close to those observed. In short, the question as to how important is distance in locational decision making in contrast to the notion that cities grow out from the core by simply adding buildings on available land closest to the growing mass, is still largely

undecided. AlphaWorld and its like could well provide insights into how such forces play themselves out.

It is unlikely that study of AlphaWorld per se would yield much insight into these questions unless a very detailed log of all decisions were available. It is much more fruitful to consider setting up a new urban world for purposes of study where the decision-making process is known in advance and where distance may be a factor in location. This would ensure that the inventors or owners of such a world would structure the log of data in such a way as to be amenable to urban analysis and simulation. The problem would be in planning such a world to take off in the unplanned fashion that has led to the kinds of organic growth characteristic of AlphaWorld. This could happen but it also suggests that many such worlds should be examined, for doubtless it will be possible to classify them into different types. The kinds of model that could be used to simulate these worlds are those based on local action such as in cellular automata (CA). CA are able to simulate growth and diffusion under very different sets of rules which could be adapted easily to the rules that are present in worlds like AlphaWorld. For example, the morphology of AlphaWorld looks like crystal growth around a core or seed, and there is a class of such models which are able to generate similar structures where the forces which determine where unoccupied cells or vacant land 'crystallises' into building can be tuned to reflect the tradeoff of agglomeration and dispersion. In fact, CA models ignore global distance effects, generating growth in such a way that global pattern is always a result of local action (Batty, 1997b).

Notwithstanding the design of the world and its placement on a server which can log all interactions, the models we might use to explore this world exist themselves in a digital environment. It is quite possible that the model of the virtual world is itself developed as a virtual environment, in the manner of the retail location model (Exemplar II). This poses an interesting regress: to use VR to understand VR and of course one might imagine the regress being pursued to several orders, or more likely different aspects of the world and its model might be explored to different orders of regress. At present, we have the design of such an urban world under scrutiny with the obvious possibility that all kinds of research into this urban cyberspace might result, from its conventional urban geography, to fully-fledged mathematical simulation, and even to its cyber-politics.

8. Next Steps

If we examine the classification we originally introduced in Figure 1, there are areas where we have no experience of virtual environments but which appear relevant to the wider study of urban applications. In this paper, we have largely ignored immersive worlds and worlds which are populated by single users, instead concentrating on environments which are much more open to many users linked together in the flesh or over the net or both. Furthermore, we have only barely touched on fictional worlds which have enormous potential for the development of theory. Theory in the form of hypothetical models always contain elements of fiction, and this is often the domain of single individuals. Virtual laboratories which bring resources to theory development in a remote context across the net or virtual field courses through which individuals learn about data, its interpretation, and its simulation are developments which we have not sketched here but which must clearly be considered as examples of virtual urban environments.

In terms of the exemplars illustrated here, each one is under continual development. The entire area is under very rapid development at present and many of the ideas, particularly those of linking different media and of displaying this media remotely across the web, have only become possible in the last few months. In terms of our four modelling themes based on navigation, decision, analysis, and manipulation, then so far there has been very little work on the kinds of navigation and decision which multiple users require when interacting within such environments. This requires very explicit and targeted research, somewhat different from the kind of applied research which is necessary to construct such environments. Throughout computing, such research has always been low key and somewhat small-scale. The single user interface which has dominated computation to date has always been left in the hands of the users and rarely do software designers tailor their designs to questions of human performance and problem-solving behaviour. Now that computation is beginning to involve multiple users, such research into the human-computer interface takes on a new urgency and it is through the kinds of virtual urban environments introduced in this paper, that progress in this domain will be made.

9. References

Batty, M. (1997a) Digital Planning: Preparing for a Fully Wired World, in **Computers in Urban Planning and Urban Management: Keynote Papers**, edited by P. K. Sikdar, S. L. Dhingra, and K. V. Krishna Rao, Narosa Publishing House, Delhi, India, pp. 13-30

Batty, M. (1997b) Growing Cities, Centre for Advanced Spatial Analysis, University College London

Carver, S. (1997) Using (Geo) Graphical Environments on the World Wide Web to Improve Participation in Social Science Research, in A. Mumford (Editor) **Graphics, Visualisation, and the Social Sciences**, Technical Report 33, Advisory Group on Computer Graphics, Loughborough, UK, pp. 11-14; see <http://www.geog.leeds.ac.uk/staff/s.carver/hpcsdm.htm>

Clarke, M., and Wilson, A. G. (1983) Dynamics of Urban Spatial Structure: Progress and Problems, **Journal of Regional Science**, **21**, 1-18

DETR (1998) **Town Centres: Defining Boundaries for Statistical Monitoring**, Department of Environment, Transport and the Regions, Her Majesty's Stationery Office (HMSO), London

Dodge, M., and Smith, A. (1998) Virtual Internet Design Arenas: The Potential of Virtual Worlds for Urban Design Teaching, Centre for Advanced Spatial Analysis, University College London

Heim, M. (1997) The Art of Virtual Reality, in P. Droege (Editor) **Intelligent Environments: Spatial Aspects of the Information Revolution**, North Holland Publishing Company, Amsterdam, pp. 421-437

Kalawsky, R. (1993) **The Science of Virtual Reality and Virtual Environments**, Addison-Wesley Publishing Company, Wokingham, UK

Kitchen, R. (1998) **Cyberspace: The World in Wires**, John Wiley, Chichester, UK

Macmillan, B. (1996) Fun and Games: Serious Toys for City Modelling in a GIS Environment, in P. Longley and M. Batty (Editors) **Spatial Analysis: GIS in a Modelling Environment**, GeoInformation International, Cambridge, UK, pp. 153-165

Negroponete, N. (1995) **Being Digital**, Alfred A. Knopf, New York

Plewe, B. (1997) **GIS Online: Information, Retrieval, Mapping, and the Internet**, OnWord Press, Santa Fe, NM

Rheingold, H. (1991) **Virtual Reality**, Touchstone Books, New York

Webber, M. M. (1964) The Urban Place and Nonplace Urban Realm, in M. M. Webber et al. (Editors) **Explorations into Urban Structure**, University of Pennsylvania Press, Philadelphia, PA, pp. 79-153

Wilson, A. G., Coelho, J. D., Macgill, S. M., and Williams, H. C. W. L. (1981) **Optimization in Locational and Transport Analysis**, John Wiley, Chichester, UK