Living with a virtual city

Alan Day, Vassilis Bourdakis and Joe Robson
Centre for Advanced Studies in Architecture
University of Bath
Bath BA2 7AY
UK

Computer models of entire cities are becoming increasingly common. The uses to which these models are put are varied and include the visualisation of proposed changes, the marketing of the facilities a city has to offer and the mapping of socio-economic data. Developments in the Internet mean that city models can be widely accessed and computer hardware and software have developed to the point where it is possible to both construct and view these models on personal computers.
This paper discusses some issues relating to the construction and use of large urban models and draws upon the authors' experience of constructing the Bath computer model which remains one of the most detailed in the world.

Any model is a simplified version of an artefact or system, produced in order to illustrate and explain some aspect of the original. This is true of both an architectural model of a proposed building and a mathematical model used by the scientist to predict physical phenomena. However, these two kinds of model are very different. One is simply a record of what might be there whereas the other attempts to encapsulate the underlying processes which give rise specific characteristics or behaviour.

Models of cities are of both types. Physical models are made to describe a city's current form and models of urban processes are produced by both physical and social scientists as they attempt to explain how cities function. Attractive though physical models are, it is often the models which encapsulate underlying social and economic processes that are more interesting. Park and Burgess (1925) who studied Chicago in the early part of this century, examined the relationships between the social and spatial organisation in the city and produced generalisable models of urban growth which explained the ecology of the nineteenth and early American twentieth century city. More recently work carried out at the Martin Centre in Cambridge has resulted in a range of models dealing with various aspects of urban form (Martin & March 1972, Echenique 1994) and Hillier and his colleagues (Hillier et al 1993) have produced models, based on the morphology of the city, which predict pedestrian and vehicular flows within cities.

With the development of computer modelling it has become increasingly common for physical models to be replaced by digital models as these have a number of distinct advantages. It is fairly easy to edit a computer-based city model and so change can be illustrated much more simply than with a physical model. Also, information from a variety of sources can be overlaid onto the urban form and, as digital information can be reproduced very simply, it is possible to make multiple copies of a virtual city and distribute them on computer disk, or over the Internet.

As a replacement for physical models, computer models therefore offer an alternative way of representing urban form. They also make it possible to bring the two approaches to modelling - the physical model and the process model - closer together. When modelling an urban process, data are being manipulated mathematically using a computer. That data will almost certainly have a spatial component and thus maps are often used to output the results. When a computer model of the form of a city exists it offers a way of explaining the results of the process modelling which is potentially much richer than through a map.

In many respects this is what the computer game SimCity does. Using simplified algorithms to represent urban processes, it allows the player to build a city and then manage it as it grows and changes. However, the city is notional and the challenge for researchers is to make simulations which can accommodate the complexities of real cities and map them onto a form which is an accurate representation of the city’s streets and buildings.

Constructing Virtual Cities

The starting point for constructing the form of a virtual city is normally a commercially available computer-aided design (CAD) package capable of creating the surface geometry of the buildings, roads and landscape. One significant advantage of any computer model is that scale does not have to be decided before the model is begun although it is still necessary to make a decision about how much detail to include. Too much and the model becomes very large making its manipulation on the computer tiresomely slow, too little and the city becomes unrecognisable as one moves down to street level. There is also a decision to be made about how detail is to be represented and there are two main ways in which this can be done. The amount of geometry can be increased, or photographic texture maps can be ‘pasted’ onto the buildings in order to represent facades.

These two approaches are very different; the first involves increasingly accurate surveying of the building stock while the second requires a comprehensive photographic survey of all the buildings. If the intention is to use an urban model for design, to provide a context within which the architect can work, then it is better to have a model with reasonable amounts of geometry. If one wants a model which the tourist can use to navigate round the streets and see what a city has to offer, then it is probably better to reduce the geometry in favour of photographic texture maps as this results in a model which is easier for the computer to manipulate and includes more of the visual clues, such as shop fronts and signage, which most people associate with a particular place.

The construction a computer model of any city is a time-consuming business and is not all that different from constructing a physical model. The geometric faces which make up the objects in the model have to be created and assembled and then colours and textures are applied to these faces prior to viewing (figs 1, 2 & 3). As with a physical model, there has to be an initial source of dimensional information and this raises the issue of how accurate the completed city should be. Ordnance Survey now produce digital maps of most cities and these provide a convenient basis for urban modelling. However, they are at 1:1250, which is a small scale to use for anything other than simple building outlines. Also, they include no information about building height and shape which have to be obtained from a street-by-street survey.

One of the most comprehensive data sources for urban modelling is aerial photogrammetry where stereo pairs of aerial photographs are viewed using binocular vision, allowing the operator to digitize objects in three dimensions. The great advantage of this technique is that the whole city can be captured very quickly and the viewpoint means that roof geometry is accurately represented and inaccessible spaces behind buildings are clearly visible. However, the greater the accuracy the larger the scale of the aerial photographs has to be and these can be difficult to obtain commercially as most stereo aerial photography is carried out at an altitude of around 3000 metres making each building fairly small on the resulting negative.

Moving through the city

The construction of a virtual city is just the first step. Once it exists it has to be made accessible and the first thing the user needs to be able to do is to move around. When navigating an urban model, either at street level or over the rooftops, one wants to be presented with a recognisable facsimile of the real thing. It is one thing to create an animation sequence where frames are pre-rendered and then linked together into a ‘movie’, but quite another to give the user interactive control of a model so that they can move through it in any way they choose. With an animation sequence it does not really matter how long it took to render each frame and, with in a complex scene, it is not unusual for the rendering to take may hours. However, interactive movement requires the computer to render each frame in real time and this is not a trivial task as at least 5 views per second have to be generated for movement to appear at all realistic, and 25 frames per second are required if it is to be completely smooth. This places a heavy demand on the computer and requires the model to be structured so that rendering can be carried out as efficiently as possible.

Viewing alternative versions of the model

Interesting though it is to move through a computer model of an existing city, it is unlikely to be as fascinating as walking around the real thing. However, if one can use the model to represent alternative versions of the city then it becomes both more useful and more compelling. Using the existing city model as a template, it is possible to construct not just one model but a whole series of models showing the city at different times. Using historical information the past can be reconstructed (figs 4 & 5) and, by including proposed schemes, the future can be anticipated.

Accessing information

As one can associate any kind of computer data with objects in a model, it is possible make links between buildings and various kinds of related information. Thus clicking on a building in the model can give access to historical information, such as drawings and written descriptions, or can provided up to the minute information on its current use. All sorts of organisations and individuals are now creating information for distribution via the Internet. Much of this is available on the World Wide Web, a part of the Internet in which information is made available as ‘pages’ containing text and illustrations (fig 6). These pages are created so that words or pictures can be highlighted and then linked to other pages of related information held elsewhere on the Internet. By linking objects in the model with information held on the Internet it begins to display some of the characteristics of the real city where the facades are not just empty shells but contain information on the activities within.

Making urban processes explicit

As an accurate representation of the form of a city, a computer model provides an ideal way of mapping urban information. Over the past ten years geographical information systems (GIS) have been used increasingly to map data which has a spatial component. Many local authorities have GIS systems which they use for a range of application, from the mapping of socio-economic data to the location of archaeological sites. Although for many purposes maps are quite sufficient, there are some applications which benefit from three dimensional representation. For example, the movement airborne pollutants can now be modelled using computational fluid dynamics programs but the output can only be fully understood in 3D. Similarly, crime incident surveys or energy use surveys make more sense when mapped onto the 3D form of a building rather than a map, particularly in areas of high density.

Existing Virtual Cities

There are a large number of computer models of existing cities, many of which can be accessed on the World Wide Web. These display considerable diversity; some are little more than crude block models of a small area while others are extensive and include considerable detail. Generally, existing models can be classified as having been created for either commercial or research purposes.

Commercial models

These range from models created by architects to illustrate a current project in its urban context to ones which are used as a vehicle for advertising businesses within a specific area. An example of the former is the model of Chicago which was started by Skidmore Owings and Merrill in the 1970s and has been built up over the years as new projects have been designed in the city. By their very nature architects' models tend to have variable amounts of detail with their own buildings modelled accurately and the surroundings only shown in a simplified form.

Models designed for advertising purposes and made available on the World Wide Web are becoming increasingly common. Virtual Soma is a model of a couple of dozen blocks in San Francisco where the geometry is relatively simple with photographs of the shops pasted onto the facades. Clicking on a facade takes you to that business's home page thus giving direct access to information on the services they have to offer. Some models focus on urban change, such as that produced by ART + COM which is a virtual reality model of Berlin which concentrates on its civic buildings and allows the viewer to switch between alternative schemes for specific sites. Others, such as those offered by BigBook are adjucts to their yellow pages service where 3D 'maps' of cities are produced in order to enhance the telephone directory by making is possible to identify companies in a specific area.

Research Models

One of the first comprehensive urban models in the UK was that constructed of Glasgow by ABACUS at Strathclyde University Most buildings were represented as simple box outline although the major monuments were modelled in more detail. More recently ABACUS has constructed a model of Edinburgh Old Town (Grant 1993) which has been used by the Old Town Renewal Trust to promote conservation strategies and to illustrate development opportunities in Edinburgh's medieval core. The Center for Landscape Research at the University of Toronto has developed a set of virtual reality tools to facilitate landscape design and these have been used in a range of urban contexts including Pittsburgh, Ottawa and Montreal (McCullough & Hoinkes 1995 and Danahy & Hoinkes 1995). Current projects which are accessible via the World Wide Web include Virtual Los Angeles (fig 7), which attempts to link CAD and GIS, a 3D map of Santa Barbara where US census data is mapped onto the urban area; CitySpace, a project where children from around the worldshare stories, pictures and sound and create a 3D navigable city model, and DIVE a multinational European project to link virtual reality with data mapping.

Development of the Bath Computer Model

For the past four years a detailed three-dimensional computer model of the city of Bath has been developed in the Centre for Advanced Studies in Architecture (CASA) at Bath University and is currently being used in a number of ways in the city (figs 8, 9 & 10). The model was created in order to visualise change in the city and that determined how it was originally constructed. Most of the data for the model came from aerial photogrammetry which was originally used to construct both the building geometry and the landform of the Avon valley. The model was constructed on personal computers running AutoCAD release 12 and 13 in order to ensure that it would be compatible with software used by practitioners.

Throughout the project commercially available software has been used wherever possible with in-house applications being limited to program customisation and software to link other applications together. Blocks (instances) were used extensively to reduce the size of drawings and improve the structuring of repetitive information as well as to facilitate the creation of databases of building elements such as windows, doors and columns. A fairly extensive and strict set of rules was devised in terms of layering, colouring, and general structure of the individual drawing files. Surrounding the model lies a 10km x 10km grid following the contours of the countryside around Bath onto which has been ‘pasted’ a texture map derived from Ordnance Survey 1:25,000 maps which were scanned into a computer and then touched up in a paint program in order to look realistic. This means that views from the city’s historic core include the backdrop of green hills which is so important to Bath’s setting.

The Bath Model is split into approximately 160 sub-models, each cross referenced with respect to the Ordnance Survey National Grid. Each of these sub-models is around half a megabyte in size so as to allow them to be copied onto a floppy disk and given to an architect to be used as the context for a new or modified building. The finalised scheme can then be returned to the total model for viewing. Initially this could only be done either within AutoCAD or by using rendering and animation software, such as 3D Studio, which although capable of producing high quality images, was non-interactive. A viewing position or animation path had to be determined and then the computer left to complete the rendering which could take may hours. In order to make the model truly interactive, to allow users to navigate thought it and decide where they want to look as they go along an alternative approach had to be considered.

An initial strategy involved pre-rendering multiple view of the city and linking there to multimedia software (Day 1994) but this was never fully implemented. Recent developments in the Internet mean new possibilities are opening up and these are currently being explored. The World Wide Web provides a standard graphical format for navigating the Internet where pages of information can contain text and pictures and which can also be linked to other media, such as sound clips, animations files and video. In 1995 an extension to the Web appeared which allowed three dimensional computer models to be added. This was an extremely important development because it offered a way of sharing 3D models across the Internet, allowing the models to be viewed and manipulated on any computer. The system is based on the Virtual Reality Mark-up Language (VRML) which defines how models should be structured so that they can be viewed by ‘browsing’ software which resides on the remote computer.

The system works in the following way. The 3D model can be constructed using any CAD package capable of exporting data using the DXF format. This model is then converted into VRML and held on a host computer. Any other computer, anywhere in the world, which has access to the World Wide Web can contact the host computer and access the model. The model is then copied across the Internet to the remote computer and the browsing software uses the processing power of that computer to manipulate the model. There are a number of significant advantages to this system. The browsing software is free and can be made available from the site which holds the model. So, provided a user can access the Web, they can download the browsing software before they view the model. There are a whole range of browsers available, for PCs, Apple Macintoshes and Unix machines, so there are no constraints on the kind of computer that can view a VRML model. Of course, the smoothness of interactive movement will be a function of the power of the computer one is using.

In-house software has been written by CASA which converts the AutoCAD model of Bath into VRML. Interactive movement within VRML means that the computer has to process a great deal of information quickly and if the entire Bath model were input in full detail then, even on a powerful machine, movement would be very slow. One therefore has to employ strategies to reduce the amount of processing needed without making the model too simplistic. One of these is to construct each element in the model at a number of levels of detail and to choose which level of detail to view depending on the camera position. Thus objects close to the camera are seen in full detail while those further away are seen in low resolution. The urban block around the Abbey which includes the Roman Baths and Pump Rooms was originally modelled using approximately 15,000 polygons. However, this can be reduced to under 100 polygons which is sufficient when viewed from afar. The Bath model has been organised in four levels of detail.

Level 1 a simple volumetric description of each terrace with a flat roof at the average height for that terrace.

Level 2 each building is modelled with accurate wall and roof geometry and tagged as a separate object in the model. This means that each property in the city can be identified and used for data linking in VRML.

Level 3 windows, doors, parapets, party walls and freestanding garden walls are added.

Level 4 architectural detail such as chimney pots, string courses and pilasters are added. At this level some photographic texture maps are also included for windows and shopfronts.

When the model is accessed, level of detail 1 is loaded along with the geometry of all the roads and landscaping. The additional levels are then loaded and unloaded depending on the camera position. On a Silicon Graphics workstation one can navigate through the city reasonably smoothly with only slight pauses when a lot of extra detail is loaded. On a Pentium-based personal computer movement is slower, but still acceptable. One particular facility that has been added to VRML is the ability to store a number of different versions of a building within the model and switch between them in real-time. This is particularly useful for comparing the impact of alternative proposals for a site from anywhere in the model.

Using the Model

The Bath model has been used in a variety of ways since it was originally constructed. To date, development control has been the main use with a number of schemes being consider. These are normally at the instigation of the local authority who recommend that schemes are modelled in order to facilitate discussions during the design phase and for presentation to the planning committee. A number of different approaches are used depending on the circumstances. For example, a school located on the slopes in front of Lansdown Crescent required a new sports hall and arts block and these were modelled in the computer. Long distance views of Lansdown Crescent were taken from the other side of the Avon valley because of concern by the planners that the sports hall might be intrusive. The study demonstrated that this was not the case and that close-up views from adjoining houses were more critical. Animation sequences were produced to show the proposed organisation of the whole site and photomontage views created to show the implications of alternative schemes from particular viewpoints. These were used during discussions with the planners and their objectivity was helpful in that the planners felt that what they were seeing was an accurate representation of what they would get and even the architects were sometimes surprised at the full implications of what they were proposing.

In that example all the modelling work was carried out in CASA, but in a more recent scheme for the redevelopment of the Southgate shopping complex in the city centre the relevant area of the computer model has been given to the architects so that they can use it while they are designing. The final scheme (or schemes) will be imported into the city model for viewing during discussions with the local authority and presentation to local residents. In another example, the development of the Royal United Hospital, the model is being used to consider the disposition of new buildings and the design of a new landscaping scheme. In all these cases the cost of using the model has fallen on the client seeking planning permission rather than on the local authority who simply recommend its use as a way of assisting the decision-making process.

In addition to its use in development control, the model has also been used to widen the public debate about the city should develop in the future. As part of the Bath Fesival there is an annual ‘Drawing on Bath’ event in which local people are invited to make proposals about how they would like to see problem sites in the city developed. Many of the resulting schemes have been included in the city model and this has proved helpful in raising public awareness of the issues and illustrating that not all new developments need not all be designed as a pastiche of the Georgian original.

Conclusions

Computer technology is just getting to the point where comprehensive urban models are feasible. Hardware and software has developed sufficiently to enable the construction and viewing of these models using personal computers and to distribute them ocer the Internet. VRML makes interactive navigation accesssible to anyone using the World Wide Web without having to purchase additional software and data can be related to these models and then accessed through objects within the model.

What is emerging is something that can be used as a central resource to both broaden and deepen the debate about urban futures. Most people find traditional ways of representing architectural schemes difficult to understand. Architects are trained to be able to synthesise the three-dimensional implications of a building from 2D plans sections and elevations. Like many skills, once learned it becomes second nature and it is easy to assume that it is something which most people have. However, they do not and many people find it extremely difficult to visualise a new building from drawings. Perspective views and physical models help but the views are usually selected to show a scheme to best advantage and the small scale of most physical models makes it difficult see the implications from ground level. When computer models are viewed interactively the control of the viewing is in the hands of the viewer rather than the designer and thus a more comprehensive and objective view of what is proposed can be realised. Although there are still deficiencies, computer models are getting more realistic and more interactive all the time and the availability of city models means that the context for new buildings can be fully represented thus providing a much more comprehensive and accurate picture. The other significant advantage of the technology is that the viewing of a proposed scheme is not restricted to a visit to the planning department but can be made available at a number of locations, either in community buildings or even individual homes.

As well as using the model to visualise change it can also become a way of outputing the results of programs which model urban processes. Thus, traffic models which predict vehicle flows can be used to place the required density of vehicles in the model and these vehicles can be represented abstractly, or as full 3D objects if, for example, the implications of a proposed traffic scheme needs to be visualised at a particular location. The open structure of the model means that the outputs from existing process modelling software can be incorporated and thus the model is not a closed system but more of a canvas upon which alternative futures can be sketched. It is this notion of being able to consider the implications of possible futures relatively easily and making them widely available that is a central attraction of this kind of urban model. Their use will widen and deepen the debate about how our cities move forward into the next millennium.

References

Danahy J, & Hoinkes R, (1995) Polytrim: collaborative setting for environmental design, in The Global Design Studio: Proceedings of the sixth international conference on computer-aided architectural design futures, pp 647-658 (eds Tan M, and Teh R) School of Architecture, National University of Singapore.

Day, A. K., (1994) From Map to Model: the Development of an Urban Information System, Design Studies, Vol 15, No.3, pp 366-384

Echenique M, (1994) Urban and regional studies at the Martin Centre: its origins, its present, its future, Environment and Planning B: Planning and Design, vol 21, pp 517-533.

Grant P M, (1993) Urban GIS -the application of information technologies to urban management, in Informing Technologies for Construction, Civil Engineering and Transport, (eds Powell J A and Day R) Brunel University, Uxbridge.

Hillier B, Penn A, Hanson J, Grajewski T, Xu J (1994) Natural movement: or, configuration and attraction in urban pedestrian movement, Environment and Planning B: Planning and Design, vol 20, pp 29-66.

Martin L, and March L (Eds) (1972) Urban Space and Structures, Cambridge University Press, Cambridge.

McCullough M, & Hoinkes R, (1995) Dynamic data sets for collaboration in urban design, in The Global Design Studio: Proceedings of the sixth international conference on computer-aided architectural design futures, pp 709-718 (eds Tan M, and Teh R) School of Architecture, National University of Singapore.

Park R E, & Burgess R W, (1925) The City, University of Chicago Press.

World Wide Web sites referred to in the text:

Virtual Soma
http://www.hyperion.com/planet9/vrsoma.htm

Berlin
http://www.artcom.de/projects/stpl/WWWpaper/CyberCity.html

BigBook
http://www.bigbook.com/

ABACUS
http://www.strath.ac.uk/Departments/Architecture/abacus.html

Centre for Landscape Research
http://www.clr.toronto.edu/clrf.html

Virtual Los Angeles
http://www.gsaup.ucla.edu/bill/LA.html

Santa Barbara
http://www.geodesigns.com/gdi_s3d.html

CASA
http://www.bath.ac.uk/Centres/CASA/

CitySpace
http://cityspace.org/