

A VRML model of Bath

[Vassilis Bourdakis](#) & [Alan Day](#)

[Centre for Advanced Studies in Architecture](#), University of Bath, U.K.

The presentation slides are also available [here](#)

Abstract: The aim of this research project is to utilise VRML models in urban planning in order to provide easy-to-use visualisation tools that will allow non-experts to understand the implications of proposed changes to their city. In this paper, issues related to the construction and use of large urban models are discussed based on the authors' experience constructing the Bath computer model. Problems faced during the creation, translation and final adaptation of the original CAD model are presented, the solutions devised are demonstrated and suggestions regarding the management of similar projects are given.

Background

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 (Day et al, 1996). The model was constructed in order to assist in making the planning and development control process more democratic by providing a means by which proposals could be visualised and alternative schemes for a site compared. The urban model provides planning officers and the public with a way of considering a wide range of view of a proposed building rather than being constrained by the views submitted by the developer.

As the model was created in order to visualise change this determined how it was originally constructed. Stereo pairs of aerial photographs were used to construct the building geometry for an area of 2.5km x 3.0km along with the landform of the Avon valley covering a total area of 10km x 10km. One of the initial concerns was how VR technologies and VRML in particular would scale and adopt in the visualisation of a whole city. There are many small VR models but not any at this scale. The database is, to our knowledge, the largest and most detailed one produced as yet; the UCLA Dept. of Architecture and Urban Design (AUD) is currently building a model of the entire Los Angeles basin covering an area in excess of 10000 square miles as part of "The Virtual World Data Server Project" but it is still under construction. Cybercity's model of Berlin is another large application of VR focusing on urban planning. However it seems that most VR urban models currently available are commercially rather than research driven (Virtual Soma, Bigbook etc.)

The model was constructed on personal computers running AUTOCAD release 12 and 13 in order to ensure compatibility with software used by practitioners. For the same reason, none of the AUTOCAD Advanced Modelling Extensions (AME) were used in the creation of the model. Instead, only 3DFACES, LINES, POLYLINES, 3DMESHES, CIRCLES and entities with thickness were employed. These are the entities fully supported in the Data Exchange Format (DXF) making translations to other CAD and VR software easier.

A stereo plotter connected to an AUTOCAD workstation provided all the plan (XY) as well as elevation (Z) information for the buildings and roof geometry. Facade surveying, regarding position and size of windows, doors cornices etc. was carried out using street level photographs. An initial issue was the level of accuracy required of the surveying and it was decided that an accuracy of about half a metre across the 3km of the model would be sufficient for most urban design purposes. This is similar to the accuracy of Ordnance Survey 1:1250 digital maps which, although not used in Bath, provide a convenient basis for the construction of 3D urban models and ensure that any plans produced from the model are consistent with the maps.

Throughout the project, commercially available software has been used wherever possible with in-house applications being limited to program customisation. A couple of packages claiming direct creation of 3D models from photos were examined, but the output looked more like an elastic stretched membrane than buildings with

accurately defined walls. Such packages are more suitable for landscape design or for relatively small object modelling rather than urban scale work. The process is also fairly laborious and results in much lower accuracy.

AUTOCAD BLOCKS (instances) were used extensively to reduce the size of drawings and improve the structuring of repetitive information as well as facilitate the creation of databases of building elements such as windows, doors, columns, trees, etc. A fairly extensive and strict set of rules was devised in terms of layering, colouring, and general structure of the drawing files.

The Bath Model is split into approximately 160 sub-models. Sub-model size range from a few hundred kilobytes up to one megabyte (as AUTOCAD drawings). Consequently, urban blocks fit on floppy disks and can be sent to local architects to be used as the context for new or modified buildings. The finalised scheme can then be imported to the total model for viewing. As the Internet reaches more architectural and engineering practices, this could be developed to an online database.

Visualisation of the Bath model was done either within AUTOCAD (for wireframes and plots) or by using 3DSTUDIO, a rendering and animation package. A combination of 3DSTUDIO, ANIMATOR, and PHOTOSHOP were used for photomontages and video presentations. In order to enhance the views from the city's historic core, the landform mesh spanning 10km x10km was "pasted" with texture map derived from Ordnance Survey 1:25,000 maps. These maps were digitised and then touched up in a paint program in order to look realistic.

Development of the VRML 2.0 model

Database Settings

The **units** used on the CAD model were metres with four decimal points and the **co-ordinate** origin was at the UK O.S. origin point at Lands End, Cornwall (Southwest tip of England). Relative to this, Bath Abbey is at 375,120 metres on X, 164,760 metres on Y and 25 metres above sea level (Z). AUTOCAD being a CAD package had no problems dealing with such numbers, but, following the exporting/translating of the model to VRML it was found that navigation was limited to "jumping" across whole urban blocks at a time since browsers use integer mathematics for navigational calculations. This led to a series of rounding errors and great problems with Z buffering. Trying to rotate or spin (in the Examiner View) the city was impossible since browsers rotate about the co-ordinate origin; in our case Cornwall. Therefore, the co-ordinate origin was moved to the middle of the city. This solved all navigation, rotation and most Z buffering problems as well as reducing the size of the model. However, if Bath is going to be linked to a more comprehensive model of the UK, this rejection of the O.S. origins must be carefully considered.

Most, if not all, CAD software use a world base co-ordinate system; X and Y for the plane definition and Z for heights. On the other hand, most VR software, including VRML, swaps the Y and Z axes, a problem which is addressed in most DXF to VRML translation software. However, a simple transformation matrix in the VRML file corrects this problem although this affects previously defined viewpoints. Excluding these from the effects of the transformation matrix makes the file difficult, if not impossible, to comprehend and edit manually. Failing to exchange the Y and Z results in models that cannot be "walked" through, since the browser-perceived walking is carried in X and Z meaning the visitor comes flying from the sky to the ground. Similarly, gravity only works if Z is properly oriented.

File Structure

The Bath database is very large, close to 90 megabytes in 3DSTUDIO binary format. The initial CAD drawings had approximately 3.5 million triangles. The VRML database triangle count is approximately the same size due to the restructuring of the CAD database and texturing. In order to allow properties to be identified, thus enabling data mapping and linking, walls spanning across a series of buildings were replaced with individual walls for each property, thus increasing the polygon count. Additionally, extra triangles were created for the shop front texturing and a low polygon representation for each urban block was built to support multiple levels of detail. However, the use of textures in building windows and billboards for trees helped by reducing the polygon count dramatically.

Node types used in the VRML model are mainly INDEXFACESETS for the AUTOCAD 3DFACES, POLYLINES with thickness, CIRCLES with thickness, POLYGON MESHES and primitives such as boxes, spheres and cylinders. More than 95% of the model is constructed with INDEXFACESETS due to the nature of the original database. An attempt was made to translate CAD spheres and cylinders to VRML primitives whenever possible, since a CAD sphere is expressed as 64, 128 or even more triangles; hardly an efficient representation. Similarly a box is defined as 12 triangles making the VRML file, bigger and difficult to read or edit manually.

The model presented in this paper is approximately one sixth of the whole city of Bath 3dStudio model; just under half a million triangles. Expecting any current computer to be able to load all that existing information is not realistic. Furthermore, the need for a useable solution on fairly low cost computers dictated the need for very careful structuring of the available information.

Initially, a classification of all building elements was attempted; i.e. typical 3 storey building, typical 2 storey, mews house, double pitched roof, corner property, 4 pitched roof, etc. Soon the list became too long and complicated to manage and therefore this idea was dropped as far as the main building components were concerned. Tables of windows, doors, columns and roof-lights as well as materials and colours were created and used extensively as VRML instances (utilising EXTERNPROTOS and facilitating editing of files).

Instancing is one of the most essential features of CAD programs; the ability to define an element which is later repeated, without having to re-specify its geometrical properties, colour etc. For example, the VRML version of the Bath Abbey was reduced from 9 megabytes to 1.5 megabytes with instancing. Simplifying the geometry even further and replacing complicated parts of the geometry (windows, cornices) with texture maps reduced the size to just 240 kilobytes. Although VRML has always supported instances, until now, there has been no translating software capable of generating instance aware VRML files from CAD ones.

Care was also taken to limit the amount of detail included in the urban model since there is a tendency for CAD operators to instance everything, even information that is hardly visible, and thus create unmanageable files. For example, instancing a fairly complex modelled chimney pot or roof-line ornament 200 times places a considerable load on the browser leading to an unusable VRML file with little visual improvement.

As mentioned earlier, the structure of the VRML model is of vital importance. An urban block of Bath modelled fully with textures and windows reaches 10,000- 15,000 triangles. One cannot expect that more than a few such blocks will be easily navigable. Instead a low polygon count representation of them should be used when the camera is afar using levels of detail (LOD). A typical urban block can be represented using 30-50 triangles, showing roughly the borders of the built block, when the camera is more than a few hundred metres away. Although a few of the currently available VRML toolkits have a LOD auto-creation feature they are only suitable for freeform shapes and landscapes and not buildings.

LOD calculations are computing intensive and there is a threshold of acceptable LOD use, versus geometry / texture use. For example, deciding to add textures on building facades, and switching them on and off per building using LOD nodes will bring the browser to a halt, not because of the burden of loading all these textures, but due to the need to do all the LOD checks for each building on each camera movement! It is better for the browser to do tests for 4 LODS per urban block than 200. This leads to a sub-structuring of the model into streets within each main urban block.

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 (typically under 200 triangles per urban block). Roads, pavements and landscape areas are also added in. All level 1 data is placed together in the loader file.

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. Description hints are set so that the name and address of the property is directly accessible.

Trees that are within the urban block are also switched on (as billboards). Typically Level 2 switches on at approximately 150 metres from the camera.

Level 3 windows, doors, parapets, party walls and free-standing garden walls are added. Windows and doors are defined as single faces "floating" in front of the walls (usually at 5cm) and instanced from EXTERNPROTOS definitions facilitating remote management. It should be noted that not all windows, doors etc. of an urban block are under one LOD node. LOD nodes are created on the basis of keeping concise, more or less square (in plan) areas together. This usually means organising them per street facade although a very long street will require more than one. Level 3 typically switches on at 90 metres.

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 shop fronts. The Level 3 structure is kept; Level 4 switches on at approximately 60 metres.

When the model is accessed, level of detail 1 is loaded together with the 10km x 10km mesh around Bath. The additional levels are then loaded and unloaded depending on the camera position. On a MAX IMPACT SILICON GRAPHICS workstation one can navigate through the city reasonably smoothly (at 5 frames per second) with only slight pauses when extra detail is loaded for the first time. On a Pentium-based personal computer movement is prohibitively slow.

The use of LOD is difficult to implement in landscapes. A 20x20 grid landscape for the 500 metre square piece of the model will have to be split at least in 4 (in plan) in order to use LOD effectively. However the edges, where the low level representation joins the high level, creates gaps and steps that cannot be corrected unless the low and high polygon count versions of the landscape have common edges, which defies the reason of having a low polygon count representation of the landscape in the first place. Tests carried with a high triangle count landscape showed a considerable slowdown and therefore this issue is still under consideration. The high polygon count of landscaping creates extra problems as it is also used for the navigation/collision detection when moving with the eye position a set distance above ground level.

Viewing considerations

VRML authoring is definitely not a What You See Is What You Get business; everything relies on the browser. There are currently over a dozen VRML 1.0 browsers for a variety of platforms, including Macs, PCs, Acorns, and UNIX workstations, and a handful of VRML 2.0 browsers, the numbers of which are growing quickly. The underlying rendering libraries that these browsers use determine to a large extent the rendering quality (or lack of). For example, the limestone coloured buildings of Bath are rendered red on one browser, brown on another and shiny white on a third. This problem is one of the most difficult to tackle and one that the VRML content creator has little or no control over. Therefore, a series of tests were carried out to make sure that the results are acceptable on all viewing environments and the results have been fed back to the browser developers.

Another such problem is in the use of textures. The graphics libraries used by most PC based browsers impose limitations on texture size; most limit each texture to 128 by 128 pixels. This is acceptable for small signs and tiled textures, but it affects building facades or shop-front textures. Furthermore, it indirectly affects bandwidth by prohibiting the use of "texture strips"; a series of textures stitched together in a long strip, loaded once and placed (targeted) in the relevant buildings. Additionally, PC graphics capabilities are, on average, quite low, forcing the browser to use software emulation for texturing which slows the process substantially. Attempting to texture map the landscape around Bath (25 2km x 2km tiles) led to serious texture swapping problems; the textures alone consumed almost 2MB of texture memory. By the time texture maps on the facades of a few urban blocks are added, the texture memory needs are only matched on very expensive workstations.

Finally, lights cause serious problems across the range of browsers available. Directional, omni, spot lights with different intensities and colours are not supported in a uniform way. Some PC browsers disregard all lights and place default lights in pre-calculated positions, others use a head-mounted light which produces very flat renderings. In many cases, browsers modify the ambient light and some of them don't support coloured lights.

One of the main aims in the construction of the Bath model is that navigation will be easy and trouble-free (Kaur et al, 1996). Early experiments showed that "walking" at ground level was very confusing and users would disorientate themselves in a matter of seconds. Collision detection and gravity did help but proved insufficient; and a side effect was slower navigation. The lack of textures and road level detail that users could identify was one reason, and the ease with which someone could move from one end of the model to the other and the lack of concept of time and effort needed for a kilometre walk was also important. Objects like moving cars and buses could help to add a sense of scale to the whole model as well as provide traffic direction hints. Camera viewing angles used are also crucial in creating correct perspective. The use of shop-front textures in the area east of the Abbey improved matters substantially. However, most of the navigation is done above roof level with better results as the model feels more like a 2D map and finding streets and buildings is much easier.

Tools Used

The first large VRML project done in CASA was the Map of the Future (London's West End) VRML1.0 model (Bourdakis 1996). Standard SGI tools (dxfToIv and ivToVRML) were used for the CAD to VRML translation. In this project the following issues were identified:

The concept of a VRML object is completely different from the way CAD programs structure data. The notion of layering, organising data, geometry etc. in groups cannot pass onto the VRML files successfully. Similarly, the three plus level of hierarchical structure in VRML models cannot be easily and intuitively emulated in CAD software forcing manual editing to set LOD's, links, descriptions etc. The differences between the constructional and visual focus of CAD vs. VRML models are apparent.

As material information is repeated for each object in a VRML file, global editing is impossible. Also, there is always something that is not translated as expected, or some information which gets lost in the process and it is quite normal to spend a few hours or even days, hand-optimising the translated file. Such optimisation results in files that are useable on slower computers and are more network friendly because of their smaller size.

Editing a VRML project is a balanced process, where it sometimes pays off to do the editing in the initial CAD package and go through the translation/optimisation stages again rather than trying to modify the VRML file directly, particularly as this is usually done using a text editor rather than visually. Even visual tools can create problems since there is no guarantee that editing a VRML file using authoring software and then saving it will result in the right structure, nodes and names.

Consequently it was decided to develop custom software for the translation of the CAD files to VRML2.0. Initially, the Bath CAD model had no LOD hierarchy in the VR sense of the word. Objects were structured on a per material basis since that is the most convenient way to assign colours and textures in 3DSTUDIO which was used for all visualisations at the time.

A strict standard was set for the reorganisation of the existing CAD drawings. The standard's development is an ongoing process and the CAD model's updating is a moderately complicated process since in many cases whole areas have to be retraced and built from scratch. Utilities facilitating the model reconstruction have been produced. As explained earlier, none of the AUTOCAD specific solid modelling extensions (AME for r.12 and ACICS for r.13) were used since these create proprietary databases that cannot be easily queried externally by custom software and create CAD files not fully compatible to the Data Exchange Format (DXF). The method devised uses AUTOCAD layers, colours and extended data to enable LOD creation from the software as well as PROTOS of materials, windows, doors etc.

The greatest problem with model translations is the structure of the geometrical description itself. In most CAD packages, the operator can define surfaces that are perceived as double-sided. That means, considering three points in 3D space, the surface defined by triangle (A, B, C) is the same as the one by triangle (A, C, B). VRML browsers, and the VRML language description (VRML2.0 ISO Spec), declares surfaces as being single sided and defined anticlockwise. Browsers have the option of rendering double-sided faces - at the penalty of a twofold speed decrease. Furthermore, the VRML language supports geometry hints, which most PC browsers discard, producing unacceptable results. When building such a large database, one cannot afford double-sided rendering

(when available) since the navigation speed be unacceptable and it pays to ensure that face normals are correctly set-up. This has the extra advantage of being able to avoid the generation of normal information in the VRML file resulting in a 20% smaller file with no drop in quality. 3DSTUDIO has a normal correcting add-on but using it on anything larger than a thousand triangle model is not worth the effort; it is better to "trace" over the existing model and create a new one.

Future Work

The model presented here is under constant development in terms of content, delivery techniques and interface design. Aiming to produce an easy to use VR visualisation tool for non-experts is a Human Computer Interface (HCI) nightmare (Charitos et al, 1996). 3D interfaces are increasingly complicated and demanding (Brelstaff, 1995) and this is an area that is currently addressed in CASA. To facilitate usability, animated cameras are included producing virtual tours of the city. Directional sound will follow, improving spatial awareness and providing an invaluable educational tool.

There are two distinct areas towards which work is carried out; the Research / Engineering and the Recreational.

On the **Research/Engineering** front the following are under consideration and/or development:

- Data mapping (i.e. pollution data, crime statistics, land use and value, energy consumption in the city, temperature, etc.)
- Integration with traffic modelling software.
- Basic editing, object switching, inserting of objects, scaling, translating rotating.
- Simple landscape manipulation, alternative season simulation, plant growth etc.
- Improved navigation features.
- Links to historical information held in an external database.

Regarding the **Recreational** use of such a model and looking forward to the next six months (when DIRECT3DTM and Pentium Pros will be considered the norm):

- Creation of a search engine for businesses, sites, addresses, etc.
- Enhancements in the navigation
- 3Dsound providing historical information on the various monuments
- Area recognition tools, extensive texture mapping, street furniture, etc.

Java is expected to play a significant role in some of these developments (dynamic data mapping, sound synthesis from small text files to be used in narration etc...)

Conclusions

The casa study described above shows that the VRML technology is suitable for the visualisation of a city 2.5km x 3.0km and well over three million polygons. Consequently, urban databases can be created using industry standard tools although the process is not short of problems. A set of construction criteria and specifications are currently being refined by CASA based on the experience gained in building the urban models of Bath, London West End and more recently Gloucester city centre. The specifications will ensure scalability and compatibility of the models produced and they will control the level of performance obtained.

Recently, there have been attempts by architects and engineers to improve on the nowadays commonly used walkthrough techniques for visualisation and planning of new and historic buildings (Fritze, 1996) using VR applications. The Bath VRML case study demonstrates such a system used by architects and engineers in order to modify and most importantly assess our cities online using centrally stored sharable databases.

References

Bourdakis, V. (1996): [From CAAD to VRML: London Case Study, The 3rd UK VRSIG Conference Full Paper Proceedings](#), De Montfort University.

Brelstaff, G. (1995): Visual Displays for Virtual Environments - A review in Proceedings of the Framework for Immersive Virtual Environments Working Group pp.27-38.

Charitos, D. and Rutherford, P. (1996): Guidelines for the design and exploration of virtual environments The 3rd UK VRSIG Conference; Full Paper Proceedings, De Montfort University.

Day, A., Bourdakis, V. and Robson, J. (1996): Living with a virtual city ARQ, Vol2 pp.84-91.

Fritze, T. and Riedel, O. (1996): Simulation of complex architectural issues in virtual environments in Virtual Reality World96 Proceedings.

Kaur, K., Maiden, N. and Sutcliffe, A. (1996): Design practice and usability problems with virtual environments in Virtual Reality World96 Proceedings.

World Wide Web sites referred to in the text:

VRML2 Spec <http://vag.vrml.org/VRML97/>

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

BigBook <http://www.bigbook.com/>

Virtual L.A. <http://www.gsaup.ucla.edu/bill/LA.html>

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