To appear in Proc. of East-West International Conference on
Multimedia, Hypermedia, and Virtual Reality, Moscow, Sept. 94.

Hooking Up 3-Space: Three-Dimensional Models as Fully-Fledged Hypermedia Documents

Keith Andrews and Michael Pichler

Institute for Information Processing and Computer Supported New Media (IICM),
Graz University of Technology,
Schießstattgasse 4a, A-8010 Graz, Austria.

{kandrews,mpichler}@iicm.tu-graz.ac.at

Abstract

This paper examines the incorporation of three-dimensional models into hypermedia systems as fully-fledged documents. Their use provides hypermedia authors with an additional, powerful means of presentation. Display, linking, navigational, and authoring aspects of 3D hypermedia documents are discussed and are illustrated with examples taken from the Harmony Viewer for the Hyper-G hypermedia information system.

1 Introduction

Hypermedia evolved from the concept of *hypertext*. Unlike the typical printed book, which is read sequentially from beginning to end, hypertext is inherently non-linear: it comprises many interlinked chunks of self-contained text. Readers are not bound to a particular sequence, but can browse through the information naturally by association, following their interests by clicking on highlighted keywords or phrases in one text to bring up another associated chunk of text.

Hypermedia is the generalisation of hypertext to include other kinds of media as well as text. Individual chunks of information are usually referred to as *documents* or *nodes*, and the connections between them as *links* or *hyperlinks* – the so-called *node-link* hypermedia model. A *source anchor* is the starting point of a hyperlink and specifies the part of a document from which an outgoing link can be activated. Typically, the user is given visual cues as to where source anchors are located in a document (for example, a highlighted phrase in a text document). A *destination anchor* is the end point of a hyperlink and determines what part of a document should be on view upon arrival at that node. [Nie90] and [BD91] are good introductory texts on hypertext and hypermedia.

Whereas traditional hypermedia systems have supported document types such as text, image, drawing, sound, and film, the incorporation of three-dimensional models as hypermedia documents has only recently begun to attract the attention of the research community. A three-dimensional model is typically described geometrically as a set of connected polygonal faces positioned in three-dimensional space (3-space). In the context of hypermedia documents, two distinct kinds of 3D model can be distinguished: a model of an *object* used as a kind of three-dimensional illustration where the model itself is manipulated (rotated, scaled, moved, etc.) and a model of a *scene* through which the user navigates (walks, flies, etc.) in an exploratory manner. An example of the former is a hypermedia motor car handbook, an example of the latter an architectural walkthrough. Hyperlinks within a 3D model may be anchored to objects or groups of polygons in the model.

Such use of 3D models in hypermedia is the topic of this paper. Following a survey of related work, we examine the display, linking, navigational, and authoring aspects of 3D hypermedia documents and conclude with a look at future work.

Some research has also been done in the area of three-dimensional overview maps to help alleviate the problem of user disorientation ("lost in hyperspace" [EH89]) in hypermedia systems. We are ourselves working on 3D information landscapes for hypermedia [AK94], similar to the landscape representation of Unix file hierarchies pioneered by FSN [TS92]. However, this work goes beyond the scope of this paper and will not be discussed further.

2 Related Work

Perhaps the most well-known work in the area of three-dimensional hypermedia is that of Card et al [RCM93, CRM91] at Xerox PARC on the Information Visualiser and 3D/Rooms. They allow users to interactively explore workspaces modeled as three-dimensional rooms. Particular data sets (nodes) are visualised in rooms of their own, and doors lead from one room to another in a manner similar to hypermedia links.

Serra et al [SCT91] discuss the use of 3D object hierarchies with attached multimedia documents. Each component in the 3D object hierarchy (part-of relationships) may be combined into a *concept node* with text, image, and video documents. Links may be made from these text, image, and video documents to other concept nodes. However, there is no support for arbitrary links from (parts of) the 3D object as such, and the 3D object hierarchy itself forms the entire extent of the hyperstructure.

Smith and Wilson [SW93] describe a prototype system based on HyperCard and Virtus Walkthrough (a 3D visualisation system), in the context of an academic departmental information system. They enabled users to interactively explore a 3D model of the department: when they approached within a certain distance of a source anchor, it automatically triggered to display a corresponding text document.

In his closing address at ECHT'92, the European Conference on Hypertext and Hypermedia, Jay Bolter presented his vision of "hypertextualised virtual reality" [Bol92], combining the immediacy of virtual reality with the representative power of hypertext. He proposed "writing on the world", immersive three-dimensional hypermedia where textual and other information is portrayed directly in the virtual model in the form of posters, stone inscriptions, and the like. Links might be represented as some kind of flying device like a magic carpet, whisking the reader to another location, as a tunnel entrance, the tunnel leading to a new location, or as a "magic window" into another world (which has the advantage of giving some kind of preview as to where one would go).

The above systems are very much hand-crafted, proof-of-concept prototypes. In Graz, we have implemented a fully-integrated 3D viewer [And92, Pic93] as part of the Harmony browser [AKS94, And93] for the Hyper-G hypermedia information system [AK94, Kap93]. Model description files are stored in the hypermedia database just like any other document. Users arriving at such a node (for example, by following a link from some other node in the hyperweb) are presented with a 3D representation of the model. They are then free to explore the model interactively at their own leisure using a variety of metaphors. Hypermedia links are anchored to individual objects or subobjects within the model and may be clicked to activate the link. Such anchors may be optionally highlighted, to gain an impression of which objects have links attached.

Figure 1 shows the Harmony 3D Scene Viewer being used to provide quick point-and-click access to satellite images: a model of the globe contains hyperlinks to corresponding images from the European Space Agency's Nimbus 7 satellite. In this example, an image of the United Kingdom and its associated text are shown.

3 Displaying 3D Models

The quality and speed of display of 3D models depends largely on the available hardware and software resources. Current display quality ranges from simple wireframe models, through shaded polygons, to real-time texture mapping. Radiosity methods have the potential to further increase realism. Real-time ray tracing methods are, at present, unachievable.

To guarantee interactivity, around 20 frames per second must be generated. The well-known trade-off between display rate and model complexity thus places an upper limit on the complexity of a model for a given hardware configuration. Techniques for polygon reduction such as visibility preprocessing and multiple levels of detail [FST92, FS93] can push this limit higher, but not eliminate it completely.