The Virtual Forest: Integrating VRML Worlds and Generative Music

Martin Dupras, B.Mus, MA
Digital Media Research Centre
Faculty of Art, Media and Design, University of the West of England, Bristol, United Kingdom
e-mail: martin.dupras@uwe.ac.uk

Nada Bates-Brkljac, B.Sc Arch, MA Urban design
Faculty of Built Environment, University of the West of England, Bristol, United Kingdom
e-mail:Nadezda.Brkljac@uwe.ac.uk

Abstract

The Virtual Forest is a VRML world representing an imaginary forest, with a real-time generated music soundtrack. This soundtrack is generated by an algorithm which responds via network messages to actions happening in the VRML world.

The VRML model uses real-life data (elevation maps, aerial and landscape photography) from a real forest to model an imaginary forest area of approximately 3 square km. The forest has a small number of animated books which when touched, trigger 'state' changes in the sound producing algorithm to alter the mood of the soundtrack. This is achieved by sending URL messages to a PHP gateway which then informs a PD (Pure Data) program to change its internal state accordingly. The PD program uses stochastic processes to generate, manipulate and recycle (or feedback) a sound stream. To further enhance the user experience of interaction with the generative process, some sounds are also immediately triggered by the books; additionally, a short generative poem is also returned to the user, as the content of the relevant page in the virtual book.

This project was realized in the context of the National Creative Technology Initiative, at UWE-Bristol.

1. Introduction

This project was instigated as part of a larger research exercise, the Foresight National Creative Technology Initiative (NCTI) between the Faculty of the Built Environment and the

Faculty of Art, Media and Design at the University of the West of England, Bristol. The project sought to investigate potential areas of collaboration and benefit from mutual research baggage and approaches.

The Virtual Forest project can be interpreted in two different but complementary ways. On the one hand, it is a virtual model of an organic environment with an accompanying generative soundtrack; one can also consider the forest as a control method (or, at the limit a musical instrument) attached to an autonomous soundscape generator.

2. Background and Prior Research

2.1 VRML

Research in the Faculty of Built Environment (FBE) have been exploring the application of 3D models of urban areas and the potential for associating 3D models with a variety of other information using multimedia tools. The interactive interface offered by WWW browsers enhances previous computer aided design (CAD) based models of buildings and landscape. Recent research has since concentrated on the freely available and open-standard VRML and on the integration of 3D VRML models and databases.

2.2 Algorithmic Music

The Digital Media Research Centre of the Faculty of Art, Media and Design specializes in the creation, management, and curation of digital media art. One focus of research is the construction of algorithmic mechanisms for the autonomous or semi-autonomous creation of time-based content. Much of the effort concentrates on the elaboration of programs and scripts that enable the passing of data between applications and interfaces (such as web interfaces and sensors.)

3. Implementation

The VRML world was created using a software package called Pavan from the firm VRML Publisher, an extension to the GIS (Geographic Information Systems) software MapInfo which generates VRML worlds from 2D map data. The Virtual Forest is an imaginary world,

and thus does not represent an existing area of forest. However, in order to achieve a certain amount of realism, the forest was built using existing data from parts of existing forests.

A terrain was first built from an elevation map obtained by combining sections of existing elevation maps. This terrain was imported in Pavan and made to fit an area of about 1 square mile (about 3km square). A collage of aerial photographs was then applied onto the terrain so as to make it look realistic when viewed from above.

Several trees and bushes were created in 3D software packages and imported into Pavan. Care was taken to achieve a reasonably compact VRML world. 3D models of trees often comprise very large number of polygons and therefore a compromise had to be attained between the complexity and number of trees. One solution was to add some coarser tree models (built from fewer polygons.)

Because the area of the VRML is quite large, it was deemed necessary to focus on certain areas which would be denser and richer; viewpoints were attached to these areas so that the viewer would be steered towards these particular areas, while still allowing them to roam freely in the forest.

To create the interface to the sound installation program, 3D models of books were then created using free available 3D modelling packages and imported in the focus areas of the forest. The books were placed so as to coincide with the viewpoints.

The sound installation was created in the PD language, a real-time object-oriented iconic programming environment. PD is a variant of the Max family of languages developed by Miller Puckette. PD is a computer music language and allows one to program their own sound algorithms using a variety of signal generators and control mechanisms.

At the core of the Virtual Forest PD program are several sound generators, using both sample playback and synthesizer techniques and processors. These are fed to a network of delays and processing units which process and alter the sound signal.

When no messages are received from the virtual world, the program is in an idle state, where sparse ambient events are triggered by stochastic processes. When messages are received, additional percussive events are triggered and sent within the audio signal. This provides a degree of interactivity by providing the user with an immediate effect. Each message also

alters the state of the program and alters the settings of the different processors. Those settings slowly restore themselves to an idle state after a short period of time.

The interface between the VRML world and the PD program is achieved by means of web-based scripts written in PHP, a server-side, cross-platform, HTML embedded scripting language. Each of the books in the forest is assigned a URL which invokes a PHP script. The script simultaneously sends a UDP message to PD (describing which book has been activated) and generates a unique poem which is displayed in a separate frame in the user's browser.

The poem is generated using simple grammar rules and a list of words and expressions. The poem is meant to represent the content of the book, and inform the user that something has happened in the system.

4. Problems

Some aspects of the development of the project have faced us with unexpected challenges, and some of our prior assumptions had to be re-examined in the course of the project.

The two main challenges had to do with the delivery of the content, and with the technical realization.

4.1 VRML problems

The VRML modelling language itself presented us with some difficulties. Although the language is very powerful and compact, it is optimized for modelling buildings and geometrically simple objects. While one can obtain a realistic model of a building using a very small number of surfaces (in the order of a few tens, perhaps) using texture mapping to give the visual detail, organic and non-geometric objects such as trees can be considerably more complex. It is a recognised problem which has been resolved in 3D CAD modelling packages by providing pre-made models of vegetation (such as bushes, shrubs and trees.) For instance, even a single flower may be more complex to model than a tall building. The number of polygons is directly proportional to the rendering speed, and thus we reached practical limits to the complexity and number of trees that we could contain in one scene earlier than anticipated.

One way that the problem was circumvented was by inserting vertically-standing pictures of tree landscapes in the backgrounds; however, whereas this may work very well in environments such as computer games, where the movement of the user is often constrained by walls, the free-roaming nature of the Virtual Forest entails that the user may sometimes get very unrealistic views.

4.2 Generative Music Problems

The audio capabilities of VRML, although adequate for many purposes, are very limited. Sounds have to be loaded in memory along with the model, and the playback mechanisms allow little else than the simple playback of sound when activated by a trigger in the scene. For this project it was essential to have considerably more flexible synthesis and control capabilities.

At the time, there was no sufficiently capable client-side solution that would have allowed the kind of processes that were needed. We consequently had to make the decision to synthesize the sound stream on the server. By nature, real-time synthesis languages have a relatively high but constant CPU usage, and therefore it was not possible to synthesize more than one stream at a time. For this reason, this version of the project is site-specific and can only accommodate one user at a time.

5. Conclusion

While the realization of the project was satisfactory considering the methods and technologies employed, we expect future derivative projects to involve a few changes in methodology.

The client-server paradigm which was used for the project, where the VRML world was browsed on a client and the soundtrack generated on a server could be alleviated entirely by generating the sound on the client machine as well; at least two new recent additions to available tools which may allow this and are being investigated as possible alternative include Phil Burk's JSyn (a client-side java synthesizer language) and Macromedia Flash 5. In both cases, the server would only have to send sound material on demand, providing true web delivery with a unique soundtrack experience to every user.