

PGPGrid

Pepper's Ghost Productions Grid

Introduction

In the commercial world an increasing number of projects are being undertaken by consortia of companies who come together for the lifetime of a contract and then break up on successful completion of the work.

This form of working is commonplace in, for instance, the construction and aerospace industries, but also in the media industries (film, television etc). Such examples of collaborative working are called Virtual Organisations (VOs) in the computer science and Grid communities.

In addition to providing seamless access to a pool of computing and data resources, the Grid has been widely identified as providing a core platform for the establishment of VOs for industry and commerce.

PGPGrid

This project proposes to explore the use of the Grid for this purpose.

It brings together three complimentary but very different organisations: Pepper's Ghost Ltd (a computer generated animation company), the 3D-Matic Lab at The University of Glasgow (part of the Imaging Faraday Partnership and expert in 3D processing technology) and EPCC at The University of Edinburgh (Europe's leading high performance computing centre).

The NeSC is represented in this proposal by both EPCC and the 3D-Matic Lab.



The PGPGrid project involves the parallelisation of the process of extracting range data from an experimental 3D tv studio.

3D-TV

This studio uses 24 video cameras focused on to the subject. These cameras each take 25 pictures per second. Three dimensional information can be obtained from the parallax effect (movement against the background) observed from similarly placed cameras.

The cameras are organised in groups of 3 with 2 monochrome and one colour camera per group. Only the information from the monochrome cameras is currently used for ranging.

Generating Data

These cameras have a resolution of 640x480 which gives a data rate per pair of monochrome cameras of 14.6 MB per second. A single pair of frames has a size of 614400 bytes.

Generating Data

continued

In addition there is a colour image which, assuming we have 16 bit colour pixels, will be the same size as the pair of monochrome ones.

The colour image is not used for ranging but the colour images from all 8 groups of cameras must be adjusted once ranging is complete to blend the brightness of adjacent pixels from distinct views.

Thus about 245MB of input data is supplied per second.

Each pair of frames is processed by a java program which extracts disparities from the pair.

A disparity map is an image the same size as the original with 3 data planes, each make up of 32 bit floats, and encoding the x and y image disparities along with the local crosscorrelation coefficient of the images at that disparity in the third plane.

This disparity map has therefore a size of 3.6MB.

From the disparity map a range map, containing x,y,z coordinates in meters is built, this has a size of 1.2MB.

The process of extracting range data from one pair thus takes about 600Kb of data and generates about 4.8 MB of data.

All 8 pairs of cameras ultimately produce of the order of 900 MB of range information per second of filming, of this only 245MB have to be retained to the next phase.

Following the extraction of the range data, the range data from each of the 8 pairs of cameras has to be combined to form a .wrl model for the whole scene at that time instant.

Finally the sequential models have to be combined into a single time/space model. This stage need not be parallelised. The approximate size of the generated .wrl files is about 80MB per second of filming.

Suggested Organisation

The optimal organisation of these tasks depends crucially on the amount of disk buffer store available to handle the intermediate data. We will probably use the following approach, in which resources will be scheduled in blocks of 8 processors:

Each processor ranges 8 pairs of range files, giving a total of 64 range maps, or 78MB of range data.

Each processor in the group then takes 8 range maps and produces 1 .wrl file.

The range files are then deleted

The process is iterated for the next 64 range files.

The approach provides a compromise in terms of storage capacity at EPCC. If we used 40 procesors it will require a working space for input and intermediate files of about 1.2 GB per major loop iteration

Further Information

Project Members:
Dr Paul Cockshott

Pepper's Ghost Productions homepage:
<http://www.peppersghost.com>

PGPGrid further details can be found at:
http://www.nesc.ac.uk/action/projects/project_action.cfm?title=125