Ibero-American Symposium on Computer Graphics - SIACG (2006)
P. Brunet, N. Correia, and G. Baranoski (Editors)

Grid Architecture for Distributed Rendering

J. A. Fernández-Sorribes, C. González-Morcillo, L. Jiménez-Linares

Escuela Superior de Informática, Ciudad Real
University of Castilla-La Mancha, Spain

Abstract

Rendering is the process by means of which a raster 2D image can be obtained from the definition of a 3D scene. This process is computationally intensive and requires a lot of time to be done when the source scene has certain complexity or when high-quality realistic images are required.
YAFRID (Yet Another Free Render grID) is a system that takes advantage of the characteristics of computational grids by distributing the rendering of a scene among a large number of heterogeneous computers connected to the Internet. With that kind of systems, the time a scene takes to be rendered is drastically decreased because the parts in which the complete work has been divided (called workunits) can be processed in parallel and finally joined using an interpolation function to obtain the image result.
As we will discuss on the final section of this article, the selection of the workunit size is a key step in this class of work division. With this approach, local optimizations in each workunit could be made to obtain a better rendering time per frame.

Categories and Subject Descriptors (according to ACM CCS): I.3.2 [Computer Graphics]: Distributed/network graphics I.3.3 [Computer Graphics]: Parallel processing I.3.7 [Computer Graphics]: Raytracing

1. Introduction

The process of creating a 3D animation comprises several phases such as character design and modeling, setting textures and materials, construction of the bones and subsequent animation, lighting, and, finally, rendering.

The latter consists basically in generating a 2D image from the abstract description of the geometry of a scene plus the definition of lights, cameras, and materials. This usually is the most computationally intensive phase of the whole process and, as a result, it takes a long time to be done. That situation is even worse when the scene to be rendered is complex or when high-quality realistic images are required. Rendering is often considered to be a bottleneck in that kind of projects and even obtaining just one image could need a lot of rendering time.

In spite of the fact that the huge time dedicated to render a scene represents an important problem, rendering is also a highly parallelizable activity because each frame of the whole animation can be calculated independently of the others. In the case of static scenes, an image could be divided into fragments to be processed separately.

In order to solve such a key problem in animations, and taking into account the advantage mentioned, several approaches has been developed. Today, the most popular one is the well-known **render farm**. A render farm is a computer cluster owned by an organization where each frame of an animation is independently calculated by a single processor.

Over the last few years, 3D animation companies have been discovering a new source of benefits in the Internet. Besides using its render farms in its own productions, some of this companies offer render services via Internet. An user of these render services can make use of the dedicated cluster that the company owns. The main difference between them and the project which this document is about (apart from the frame-level division and interpolation schemes) is that Yafrid allows anyone to be part of the grid by sharing CPU cycles via the Internet.

This technique and other similar ones work by dividing a complex task (rendering a complete animation or scene)