Compressing RIB Sequences using diff

Ian Stephenson

National Centre for Computer Animation
Bournemouth University
istephen@bournemouth.ac.uk

1 Introduction

During the production of animation sequences most rendering systems store each frame in a metafile. These files may be stored on disk for some time before being processed by a render-farm, consuming considerable amounts of storage space. Though the individual files are typically compressed, it is proposed here that significant additional compression may be obtained by considering the similarities between each frame of animation.

Within a typical shot, certain objects remain static or change in only a limited fashion, while others change more radically. Only those elements which change need be recorded. This can be exploited within an existing production pipeline, and hence storage requirements reduced.

Though we will specifically consider the RenderMan ASCII RIB format, the results may also be applied to Mantra's ifd files, Mental Ray's MI files or the Binary RIB file format.

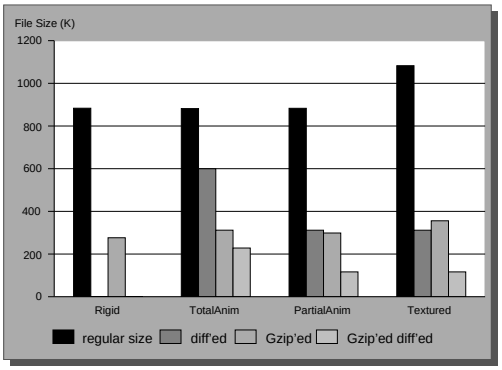

Figure 1: Compressed File Sizes

2 Generating Diff's

RIB files are simply ASCII text containing commands to specify the camera setup and geometry of the scene. To compress a sequence of these files we need to derive a near minimal set of changes that must be made to the text describing one frame to convert it to a second. While this is a non-trivial task, it has been well researched, and in fact all Unix machines are shipped with a utility to perform such a calculation: "diff".

The generation of differences (diff's) requires CPU time comparable with other compression methods, but diff's can be expanded, with no significant memory or CPU overhead. This decompression may be integrated into an alfred script allowing the compressed files to be expanded by the render farm as required.

We chose to construct diff's with respect to a single reference frame. Though the results may have been improved slightly by using multiple reference frames, or by defining each frame in terms of its predecessor the complexity, and fragility this would have brought to the system was deemed unacceptable.

3 Results

For test purposes, a simple scene was built, containing two relatively detailed polygonal models. These models where then animated in a number of ways, and the average size of the per frame differences is shown in figure 1.

The scene was first animated by moving the camera, and applying transformations such as rotation, and translation to the objects.In such a simple case, the diff's from one frame to the next are negligibly small. Though the models may be complex, they are stored only once in the reference frame, and subsequent frames record only the transformation commands.

In a practical situation some objects would be deforming. To simulate the worst case scenario, every point was displaced randomly at each frame. Even in this case compression of 30% was achieved, as the topology of the objects remains unchanged.

To represent a more typical case 50% of the points in each object were animated, and 50% left unmoved. Diff successfully separated the two sets of points, resulting in compression of approximately two thirds.

Though diff compression fairs less well when objects are deforming, it is these cases in which texture coordinates or reference geometry is often used. Adding texture coordinates to the scene significantly increased the size of the standard file, but had virtually no effect on the diff's.

For comparison, the scene was also compressed with gzip. For the more realistic test cases, gzip and diff compression produced similar results. More significantly the two methods were found to be orthogonal – diff's could be further compressed using gzip, resulting in a file typically one eighth of its original size.

When applied to RIB files from production sequences, the results were found to be variable, but broadly in line with those produced by these test cases.

[0]Presented at the PIXAR User Group Meeting, San Antonio, Siggraph 2002

[0]Supporting Code, and scripts are available from the DCT website: www.dctsystems.frerserve.co.uk