DISTRIBUTED VIDEO ENCODING OVER A PEER-TO-PEER NETWORK

Daniel Hughes, James Walkerdine

Key words to describe the work: Peer-to-Peer Applications, Distributed Computation, Video Encoding

Key Results: An extensible video encoder, which uses a lightweight peer-to-peer grid to leverage the processing power of regular PCs for the purpose of distributing the encoding of highly compressed video, for example MPEG4 and H.264.

How does the work advance the state-of-the-art?: Current video encoding technologies tend to focus on single machine solutions, while little or no work on distributed video encoding systems has been undertaken. Current work on distributed computation over peer-to-peer networks primarily focuses upon systems with heavily centralised control [1] [2]. The Distributed Video Encoder is a novel example of fully decentralized ad-hoc distributed computation.

Motivation (problems addressed): Video encoding is highly processor intensive. While the processing power available on standard PCs is continually increasing, so is the complexity of modern audio/video codecs. The Distributed Video Encoder uses a peer-to-peer network as a substrate for sharing the processing power of regular PCs for the purpose of encoding video. This allows nodes the easy construction of video encoding 'farms' from cheap standard PCs.

1. Introduction

Video encoding is a lengthy, CPU intensive task, involving the conversion of video media from one format to another. Furthermore, and of particular interest to distributed processing, input video files can be easily broken down into work-units. These factors make the distribution of video encoding processes viable.

The majority of research on distributed processing has focused on Grid technologies, however, in practicality, Grid services tend to be inflexible and offer poor support for ad-hoc interaction. Consequently it can be difficult for such technologies to efficiently exploit the growing pool of resources available around the edge of the Internet on home PCs. It is these resources that the DVE uses to provide a distributed video encoding service.

2. Implementation

In order to expedite the development of the distributed video encoder, it is being developed as a (Java) plug-in for Lancaster's P2P Application Framework [3]. The P2P Application Framework provides an abstract layer specifically geared to rapid P2P application development, reducing the burden upon the developer to understand the underlying P2P technologies and allowing them to focus on application development. This is accomplished using a set of generic application orientated services.

The framework operates over a semi-centralised network structure, which makes use of indexing peers. Specifically the DVE makes use of the framework's search, file-sharing and resource-awareness services. The Java Native Interface (JNI) is used together with the Microsoft Windows Media Encoder SDK [4] to implement the video encoding functionality. This is shown in Figure 1.

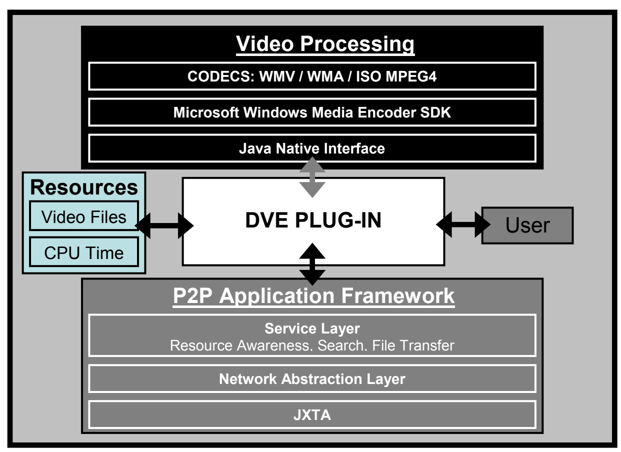

Figure 1 – DVE Architecture

Currently the DVE supports only MPEG input and is capable of encoding to ISO-MPEG4 and Windows Media Video (WMV) versions 7, 8 and 9. Support for additional video formats will be added in due course.

The encoding process used within the DVE is broken down into four stages:

1. **Peer discovery** - the distributing node uses the framework to discover what encoding resources are available on the network.
2. **Generation of work units** - the source video file is split into x chunks, where x is the number

of nodes on the network willing to process work units.

3. **Distribution and encoding** - encoding specifications are sent to each participant node, which downloads its work unit and encodes it.
4. **Recombination** - encoded work units are returned to the distributing nodes and re-assembled into a complete encoded video. If work units are not returned due to node failure, they are resubmitted to the network for encoding. This process repeats until all work units are received by the distributing node and reassembled.

3. Evaluation

A first release has been made available and partially evaluated. In a LAN environment, The DVE successfully uses the redundant computing power of regular workstations to provide a powerful distributed video-encoding platform. The graph below shows the system's performance on a 100mbps local area network using between 1 and 45 nodes (of equal speed).

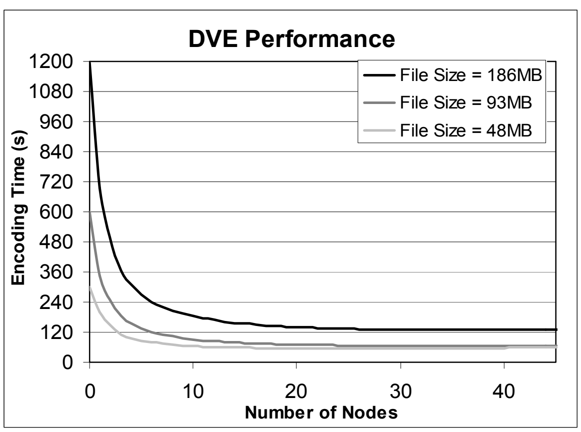

Figure 2 – DVE Performance

Figure 2 illustrates the performance of the distributed video encoder using a variety of file and network sizes. The test file used contained MPEG1 video of dimensions 192 * 144 and with 64kbps MP3 audio. This was encoded into 70kbps WMV9 format. Large performance increases over single node encoding can be observed in each case.

4. Summary and Future Work

This abstract paper presents ongoing work on the development of the Distributed Video Encoder, a means to utilise the spare computational resources of standard PCs within a P2P network. An overview of the system's design has been presented along with an initial evaluation into its effectiveness. This evaluation shows that within certain circumstances the DVE can successfully reduce encoding time. Development of the DVE is ongoing and there are a number of open issues we intend to further examine. These include:

Error Recovery: Currently, the system decides that a work unit has been lost due to node failure based on a timeout computed using the time taken for the first work-unit to be returned. In situations where there is a large disparity between the processor-speed of nodes, this approach is naïve and may result in fast nodes sitting idle before the work-unit is resubmitted. This problem could be tackled by providing the distributing node with increased awareness of encoding progress across the network.

Support for Additional Codecs: The framework currently only MPEG input and WMV / MPEG4 output. Support for additional technologies, especially MPEG-II video, the encoding of which is more processor intensive than the two currently supported output technologies would increase the usefulness of the system as would the ability to use AVI files as input.

Security: One of the key aims during the development of the DVE was to provide distributed video processing facilities to home users. The type of material home users would typically want to encode is by its very nature personal. Currently work-units are distributed as raw mpeg video chunks. Some level of digital rights management would be advantageous: allowing users to encode personal video with the service, while not allowing unauthorised persons to view whatever chunk of the material they may be encoding.

5. References

[1] Seti@Home, http://www.setiathome.com

[2] Folding@Home, http://folding.stanford.edu

[3] Walkerdine, J., Melville, I., Sommerville, I., A Framework for P2P Application Development, Technical Report COMP-004-2004, Computing Department, Lancaster University, 2004.

[4] Microsoft. Windows Media Encoder 9 series SDK www.microsoft.com/windows/windowsmedia/9series/