

Distributed Video Encoding over a Peer-to-Peer Network

Daniel Hughes, James Walkerdine

Computing Department, Lancaster University, Lancaster, UK

d.r.hughes@lancaster.ac.uk | walkerdi@comp.lancs.ac.uk

Abstract

There are few tasks commonly undertaken by users that are as heavily processor-intensive as encoding digital video. The rising popularity of video capture devices such as digital camcorders and video editing software means that an increasing number of people are regularly using their PCs to encode large video files. While the processing power available to PCs is continually increasing, so is the complexity of modern audio/video codecs. The Distributed Video Encoder uses a semi-centralised peer-to-peer network as a substrate for sharing the processing power of regular PCs. This allows nodes to utilise the spare CPU cycles of their peers and has the potential, especially in Local Area Network environments, to significantly reduce video encoding time.

1. Introduction

This paper provides an overview of the Distributed Video Encoder (DVE). A peer-to-peer (P2P) based distributed video encoding application.

Video encoding is a lengthy CPU intensive task, involving the conversion of video media from one format to another. Furthermore, and of particular relevance for distributed processing, input video files can be easily broken down into work-units. These factors make the distribution of video encoding 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 on regular user's 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 was developed as a (Java) plug-in for Lancaster's P2P Application Framework. [1]

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.

3rd party encoding tools are used for video processing. The Java Native Interface (JNI) interfaces with the Microsoft Windows Media Encoder SDK. [2] This is shown in Figure 1.

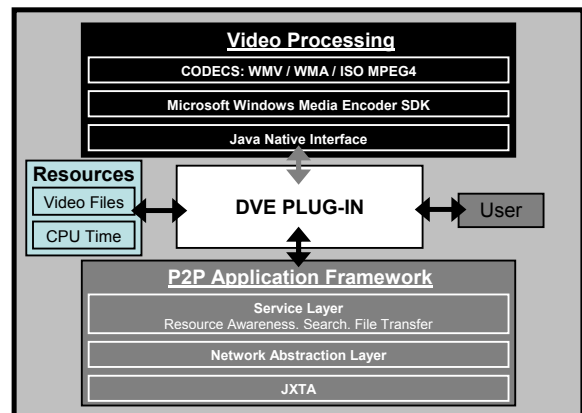


Figure 1 – DVE Architecture

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:

- *Peer discovery* - the distributing node uses the framework to discover what encoding resources are available on the network.

- *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.
- *Distribution and encoding* - encoding specifications are sent to each participant node, which downloads its work unit and encodes it.
- *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 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. The test file used contained MPEG1 video of dimensions 192 * 144 and 64kbps audio. This was encoded into 70kbps WMV9 format.

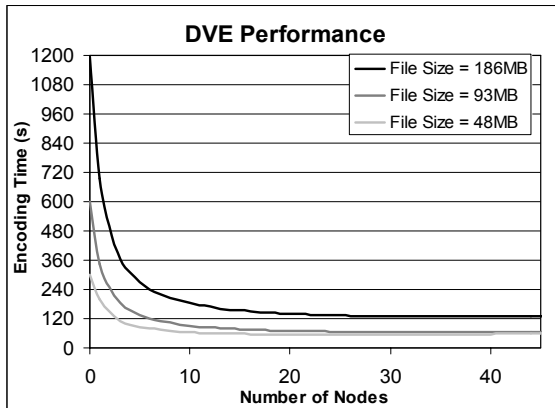


Figure 1 – DVE Performance

Figure 1 illustrates the performance of the distributed video encoder using a variety of file and network sizes. Large performance increases over single node encoding on any currently available platform can be observed.

Due to the overhead involved in splitting and reassembling the file, there are an optimal number of nodes for any given file size. Using a greater than optimum number of nodes will result in longer encoding times. The DVE will always select a number of nodes as close to the optimal as possible from the pool available on the network.

4. Summary and future work

This paper has presented our work on the Distributed Video Encoder, a means to utilise the spare CPU cycles available within a P2P network. An overview of its design has been presented along with an initial evaluation into its effectiveness. The evaluation showed that within certain circumstances the DVE significantly reduces encoding time. Development of the DVE is ongoing and there are a number of open issues we intend to examine further. These include:

- **Error Recovery:** Currently, the system decides that a work unit has been lost due to node failure based on a timeout. This timeout is 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 the nodes comprising the network, 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 supports MPEG input and WMV / MPEG4 output. Support for additional technologies, particularly 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 format input files.
- **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 tends by nature to be 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.

The DVE prototype along with the P2P Application Framework is available for download, from our departmental P2P website - <http://polo.lancs.ac.uk/p2p>

5. References

- [1] Walkerdine, J., Melville, I., Sommerville, I., A Framework for P2P Application Development, Technical Report COMP-004-2004, Computing Department, Lancaster University, 2004.
- [2] Microsoft. Windows Media Encoder 9 series SDK www.microsoft.com/windows/windowsmedia/9series/