

Parallel MPEG-1 Video Encoding[†]

*Kevin L. Gong and Lawrence A. Rowe
Computer Science Division - EECS
University of California
Berkeley, California 94720*

(Keving@CS.Berkeley.EDU, Rowe@CS.Berkeley.EDU)

Abstract

MPEG-1 is an ISO/IEC (International Organisation for Standardisation/International Electrotechnical Commission) standard for medium-quality and medium-bit rate video and audio encoding. We have developed a parallel MPEG-1 encoder that runs on a network of workstations. We have implemented a number of different motion vector search techniques. The sequential code can encode over 0.6 frames per second on an HP 9000/720 machine, and the parallel code can encode slightly less than 4 frames per second on a network of 9 of these machines connected via Ethernet.

Keywords

MPEG, MPEG-1, video compression, parallel programming

1. Introduction

The Motion Pictures Expert Group (MPEG) has produced a standard for video and audio encoding for medium-quality, medium-bit-rate compression, called MPEG-1 [3].

Our goal in developing an MPEG-1 encoder was to support the creation of the Berkeley Distributed Video on Demand System (VODS) [9, 10]. The VODS system is designed to handle hundreds of hours of video. Due to the large size and amount of video, that video must be compressed. MPEG-1 allows video to be compressed by ratios in the range of 50:1 to 100:1, depending on image sequence type and desired quality.

Commercial encoders use custom hardware and parallel encoding, which is fast and expensive. For example, the C-Cube real-time MPEG-1 video encoder uses 8 custom-designed video processor chips; a complete system is sold for \$120,000. Optibase and Optivision sell PC-based, real-time encoders based on the C-Cube video processor chips for \$20,000. Software-only encoders are slower. For example, the Stanford encoder, running on a DEC 5000, encodes at just 0.2 frames per second [2]. The Philips paral-

lel MPEG encoder runs on a POOMA machine [12]. POOMA is an experimental parallel computer system consisting of 100 nodes. Each node consists of a Motorola MC68020 CPU, MC68881 FPU, and 16 MB of memory. Later, we will compare the performance of the Philips encoder with our encoder.

We developed a parallel encoder that runs on a network of workstations to reduce the time needed for encoding. We wanted our encoder to be used by as many people as possible, including those already using the Berkeley MPEG-1 player [1]. To address this issue, we designed the encoder to be as flexible and easy-to-use as possible, allowing different methods and quality of encoding for users with different needs. Thus, our ultimate goal was to create a usable, flexible, portable, parallel encoder with reasonable quality and performance.

This paper describes the design, implementation, and performance of the Berkeley MPEG-1 video encoder. The paper is organized as follows. Section 2 describes the MPEG-1 video standard. Section 3 describes the encoder and discusses the quality and performance of the sequential encoder. Section 4 describes the parallel encoder and its performance. Section 5 presents our conclusions.

2. MPEG-1 Video

The MPEG-1 video standard is designed to encode color images at 1.5 Mb/s and to provide VHS-quality playback. The basic idea behind MPEG encoding is to exploit temporal locality. Excepting certain types of video such as music videos, the images do not change much within small time intervals. MPEG-1 video takes advantage of this by encoding an image relative to other images temporally close to it.

2.1 Data Rate

A source video is a sequence of numbered frames, $F_1 \dots F_n$. Each frame is a still image. A video player displays one frame after another, usually at a rate close to 30 frames per second (the standard television rate).

Frames are digitized in a standard RGB format, 24

[†]This research was supported by the National Science Foundation (Grant MIP 90-14940), Digital Equipment Corporation, Fujitsu Network Transmission Systems, Inc., and Hewlett-Packard Corp.