C-Cube Unveils Low-Cost MPEG Decoder

Decode-Only Chip Designed for Consumer Electronics

By Brian Case

For multimedia to become a practical mainstream PC application, the data storage and bandwidth requirements of real-time video and audio must be reduced. The usual solution to such problems is data compression, and the ISO CD 11172 MPEG (Moving Picture Experts Group) standard for video and audio compression has been widely embraced in the industry.

Starting with the demonstration of its CL950 MPEG chip about a year ago (see µPR 5/17/91, p. 8), C-Cube has been a pioneer in MPEG video compression. The 950 was a demonstration vehicle designed only for evaluation of MPEG. C-Cube's new CL450 MPEG decompression chip, developed in collaboration with Philips, fulfills their promise to make derivatives of the 950 available in volume.

Like the 950, the 450 is a decompression-only chip. MPEG compression requires much more compute power and memory than decompression, so playback only chips are a natural first step. This means, however, that the 450 is not the chip that will bring full-fledged multimedia authoring to PCs. Instead, the 450 is aimed at CD-ROM-based, VHS-quality consumer and public-use applications. Expected late this year is a compression/decompression successor to the 450 designed for PC motherboards.

The 450 and its successor implement MPEG-I, which encodes pictures in about 1/2-bit to 1 bit per pixel; the number of pixels then determines the data rate.

The 450 has about 398K transistors on a 132K mil^2 die. The maximum clock rate is 40 MHz, and it dissipates 1.5 W at this clock rate. DRAM speed required at this clock rate is 80 ns, but since the 450 is static, it can be run slower to take advantage of slower DRAM if a lower resolution is sufficient.

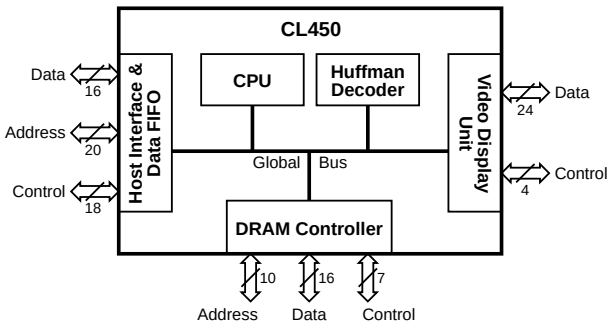

Figure 1. Block diagram of C-Cube's CL450 MPEG decoder.

Hardware Overview

The 450 has a maximum compressed data input rate (at 40 MHz) of 1.856 Mbits/s when used with the standard memory configuration. Data rates up to 3 Mbits/s can be supported, but this requires more DRAM for buffers to meet the requirements of the MPEG specification. Current CD players can deliver only about 1.2 Mbits/s of usable data, so this is the data rate that will be most commonly used.

The 450 decompresses Source Input Format (SIF) resolution bit streams. Raw SIF resolution is either 352 $\times$ 288 pixels at 25 Hz or 352 $\times$ 240 pixels at 30 Hz, but the 450 can perform real-time horizontal pixel interpolation and frame rate conversion to effectively double the frame rate and number of horizontal pixels.

The chip is packaged in a 160-pin PQFP, which provides the relatively large number of pins needed to accommodate its three separate buses: a 16-bit data/20-bit address host-microprocessor interface bus, a 16-bit data/10-bit-multiplexed address DRAM interface bus, and a 24-bit video interface bus. A rough block diagram is shown in Figure 1, while Figure 2 shows a typical CD-based system block diagram.

Consistent with its application areas, the 450 is essentially a specialized embedded control microprocessor: it has a programmable on-chip "purpose-built" RISC processor with some assist hardware and a glueless DRAM interface. The MPEG algorithm relies heavily on Huffman variable-length encoding, so the on-chip processor is augmented by a dedicated Huffman decoder.

Other time-critical functions in the decompression inner loop are assisted with special instructions, such as IDCT (inverse discrete cosine transform) primitives. By using special instructions instead of dedicated hardware, the bulk of the decompression algorithm is expressed in the processor program stored in the external DRAM. This means that the 450 can take advantage of software updates and even implement other decompression algorithms with relative ease.

Bus Interfaces

The host-microprocessor bus is tailored for direct connection to a 680x0 microprocessor, but with a few extra components it can be connected to an 80x86-class microprocessor. C-Cube plans to expand the bus of future parts to 32 bits while maintaining the 680x0 orientation. (CD-I players use 68000-family processors.)

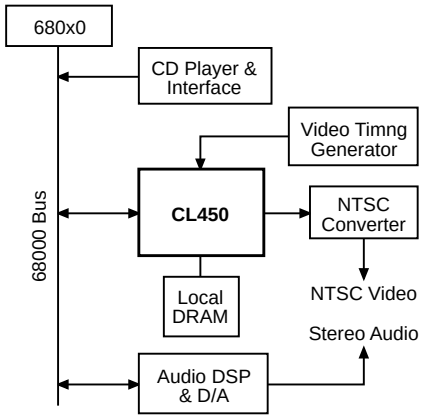

Figure 2. Block diagram of a typical CL450-based system.

Compressed MPEG data is written by the host microprocessor into the 32-byte data FIFO in either normal or DMA modes. In DMA mode, the maximum data transfer rate is 5.8 Mbytes/s. Since this transfer rate is significantly higher than the overall maximum of 3 Mbits/s decompression rate, the host processor has time left over for other processing chores, such as separating the audio and video data streams.

The host bus interface unit also includes a 33-bit counter that is used to synchronize video decompression with audio decompression.

The DRAM interface requires no glue logic to control up to 1 Mbyte of local DRAM, although only 512 Kbytes of DRAM is required in the standard configuration. Most of the DRAM is used for an input data buffer, two intermediate frames needed for decompressing succeeding frames, and the current display frame.

The local DRAM is directly byte-addressable by the host processor. This capability is used to download the decompression algorithm into the DRAM.

The DRAM is organized as two banks, and the banks can consist of 256K × 4 or 256K × 16 memory chips. With 256K × 16 organization, only one DRAM package (one bank) is needed to complete the decompression system. The DRAM controller does page-mode reads and writes in two clock cycles while random accesses take seven clocks.

The video bus provides either RGB or YUV (separate luminance and chrominance) pixels. The MPEG decompression algorithm actually computes YUV information; on-chip logic performs a simple linear transformation to generate RGB data. The original 950 required external logic for this function, but consumer applications need the higher integration of an on-chip solution.

In RGB mode, the 24-bit data bus has separate bytes for each color. In YUV mode, eight pins carry a luminance value and another eight carry chrominance. A luminance value is delivered in every clock cycle, while the red and blue chrominance values alternate (luminance has twice the resolution of chrominance). The remaining eight pins reflect the value of an internal register and so can be used as programmable outputs in YUV mode.

On-Chip Processor Architecture

The 450 is built around a proprietary RISC processor. In contrast, IIT has licensed the MIPS-X RISC processor core from Stanford for its video compression chips instead of designing its own.

The processor has 24-bit data paths throughout and the register file contains 32, 24-bit general-purpose registers with register zero always reading as zero. The instruction set is heavily influenced by the needs of the inner loop of the MPEG decompression algorithm, but it includes a fairly standard mix of integer instructions: add, subtract, shifts, and bitwise logicals. In addition to traditional instructions, C-Cube included some specialized instructions to accelerate coefficient dequantization and the computation of the IDCT.

All instructions are 32 bits long and can encode three register operands or two register operands and a 13-bit immediate. One microcode-like format allows an ALU instruction and a conditional branch to be encoded in one 32-bit instruction.

Another departure from traditional RISC architectural concepts is the block-move capability. Separate RAMs on the 450 hold the results of dequantization and IDCT for a block of pixels. Block move instructions allow the contents of these pixel memories and other on-chip RAM to be moved to and from DRAM without the bother of using a loop.

The processor is implemented with a vanilla 4-stage fetch/decode/execute/write-back pipeline, except for the high-function IDCT and dequantization instructions. For these operations, which include multiply/accumulate-like functions, the execute stage is lengthened to five stages: two stages for a 24 × 24 multiply and three stages for special logic functions related to details of the MPEG algorithm.

The Huffman decoding logic is essentially a state machine that can fetch and decode variable-length-coded data in parallel with the main processor. The RISC CPU issues a command to the Huffman decoder, and then the two work in parallel.

The decompression program is stored in local DRAM, but instructions are fetched by the CPU out of an on-chip 512-instruction buffer RAM. This RAM is not a cache—it must be managed by the program itself. The MPEG code is currently about 4K instructions, and the critical inner loops fit in the on-chip buffer.

The use of a simple but fast on-chip processor is probably a good choice for this market since it allows the chips to implement more than one compression algorithm and to be easily customized for certain markets.

MPEG Encoding Algorithm Overview

Two types of redundancy exist in full-motion video: redundancy within a single frame and redundancy between adjacent frames. MPEG compression exploits both types.

Just like an uncompressed video signal, an MPEG compressed data stream is a series of frames. Unlike uncompressed video, MPEG has three different types of frames: Intra picture (I), Predicted picture (P), and Bidirectional picture (B).

I-type frames are compressed using only information in the corresponding source frame; no reference is made to adjacent frames. I-type frames therefore yield low compression but are randomly accessible. The frequency and location of I-type frames are determined by the MPEG encoder. When random access is important, two I-type frames per second is typical.

P-type frames are compressed using information in the source frame and in the closest preceding I-type or P-type frame. P-type frames result in more compression than I-type frames.

B-type frames are compressed using information in the source frame and information in adjacent frames. These frames can be closest past or future I-type or P-type frame or *both* the closest past and future reference frames. B-type frames result in the most compression, but they are never used as reference frames so their approximation errors never propagate. The number of B-type frames between a pair of reference frames is typically zero, one, or two, but the MPEG encoder determines the number based on its memory capacity and other factors.

Motion compensation is based on blocks of 16×16 pixels. A motion vector describes how a block moves between two frames with up to one-half-pixel accuracy. The difference between the source and predicted pixels can also be included to increase the accuracy of the prediction. For P-type and B-type frames (the only ones that can take advantage of motion vectors), motion information is sent with each 16×16 block. To take advantage of the fact that adjacent motion vectors are similar (implying that adjacent 16×16 blocks tend to move similarly), the MPEG data stream encodes the difference between adjacent motion vectors. As is common in the MPEG algorithm, a variable-length Huffman code is used to encode these differential motion vectors to take advantage of the common case of nearly identical adjacent vectors.

Transform coding is also done to further exploit redundancy in blocks of error terms and raw image blocks (for I-type frames and the parts of P-type and B-type frames that are not predicted). Blocks of 8×8 pixels or 8×8 error terms (delta pixel values) are transformed from the amplitude to the frequency domain by a Discrete Cosine Transform (DCT). The frequency coefficients obtained from the DCT are quantized with the quantization for high-frequency coefficients being more coarse (because human perception of high-frequency quantization is less keen).

The resulting 8×8 matrix of quantized frequency coefficients is scanned in a zigzag manner to create a 64-element vector. The combination of DCT, quantization, and zigzag scan result in a vector with runs of zeros. So, "run-amplitude" coding is used to generate pairs of numbers: the number of zero coefficients in a run followed by the amplitude of the next non-zero coefficient. Finally, these run-amplitude pairs are encoded with a variable-length code.

Thus, decoding an MPEG data stream is a series of many simple steps: decoding variable-length Huffman tokens, disentangling matrices of quantized coefficients, inverse DCT transforms, and possibly adding pixel amplitudes from adjacent frames. Given the amount of variable-length encoding in the MPEG algorithm, the dedicated Huffman decoding logic in the 450 clearly reduces the computation required of the on-chip processor.

Although motion compensation adds a factor of three to four to the compression effectiveness, one drawback is that errors in the data stream can have effects on many decoded frames. In practice, I-type frames limit how far critical errors can propagate. With two I-type frames per second, an error would affect at most a half second of video.

C-Cube is further convinced that using specialized instructions to speed up operations like dequantization and the IDCT is better than special hardware coprocessors because it avoids duplication of logic that is already present in the CPU.

Operation Overview

The decompression process begins with the transfer of compressed MPEG data from the host to the 450. The 450 moves data from the input FIFO to the buffer section of DRAM. The Huffman decoder fetches data from the DRAM buffer and parses it into a stream of symbols (discrete cosine transform coefficients, motion vectors, etc.).

The on-chip processor dequantizes the DCT coefficients and performs an inverse DCT to convert the data back to the amplitude domain from the frequency domain. If motion vectors dictate, the amplitude data is summed with intermediate frames before complete pixels are stored in the display-buffer portion of DRAM.

Finally, pixel values are fetched by the video display unit and, if necessary, converted to RGB format. The last step before the pixels are driven on the video data bus is horizontal interpolation to double the horizontal resolution.

Applications

The 450 is not the chip that will bring full-fledged

multimedia capabilities to personal computers, but it will spark the development of a range of consumer, point-of-sale, and public-information applications. C-Cube is working with Philips, JVC, and other companies to develop products with the 450.

One of the biggest advantages offered by the 450 is a reduction in size and cost for a full-motion, random-access video machine. Previously, a full-size laserdisc player was required, but now a CD-ROM player and a few extra chips will provide equivalent capability if somewhat lower video quality. A major consumer-electronics company, for example, is developing a CD-ROM-based Karaoke machine with the 450. Informational and point-of-sale kiosks should proliferate because of the lower size and cost of the underlying hardware. In the coming years, In-Stat predicts multi-million-unit volumes for MPEG chips in PCs, CD-I (consumer interactive), and cable TV.

For C-Cube, the 450 is just the beginning. Later this year, it plans to introduce a new version for the PC market. This chip will perform MPEG encoding as well as decoding to allow PC applications to implement MPEG video authoring. With this capability, mainstream PC hardware will at last provide end users with the ability to author full-fledged multimedia presentations.

These first products implement the MPEG-I standard, which C-Cube claims is capable of exceeding VHS VCR quality and approaching laserdisc video quality. MPEG-II chips, due in late 1993, will deliver quality equivalent to current broadcast standards (at up to 10 Mbits/s). At that point, C-Cube expects MPEG-II to be embraced by the cable TV industry, become widespread in broadcast editing equipment, and become the enabling technology for digital TVs, VCRs, and camcorders. When these devices proliferate, digital multimedia could become a very important, standard part of all personal computers.

A really exciting prospect is the full-color, full-motion, digital multimedia notebook personal computer. This hypothetical device could have a 10" full-color, active-matrix LCD screen, a TV tuner, audio/video inputs and outputs, and MPEG-II all built-in to the machine. Coupled with a digital camcorder, this computer would form the core of a portable, full-function video authoring station. With a 100-MIPS CPU, lots of memory, and a large hard disk, this portable computer could perform video editing and enhancing currently possible only with studio-bound equipment.

Given that MPEG has access to the markets for VCRs, camcorders, cable and broadcast video equipment, potentially all PCs, and products in as-yet undeveloped markets, it is easy to see why semiconductor companies are excited about the profit potential of video compression chips.

Price & Availability

A sample kit for the CL450 is available immediately at $250 (the kit contains the 450 chip, a data book, object code for the MPEG program, a loader, software for initialization, and a software license agreement). Production quantities are available immediately with standard semiconductor lead times (8 to 10 weeks for an order of 10K+). In quantities of 100K or more per year, the price is less than $50.

C-Cube Microsystems, 1778 McCarthy Blvd., Milpitas, CA 95035; 408/944-6300; fax 408/944-6314.

Conclusions

C-Cube has an early start with the 450, but it will have to stay on schedule to keep from being overtaken by larger competitors. The highest-quality video compression algorithm Intel has implemented with its DVI chips is Intel's proprietary PLV. C-Cube claims PLV has inferior picture quality and that the DVI chip set requires 8 Mbits of VRAM (costing $50 to $60) while the 450 requires only 4 Mbits of DRAM (costing $10 to $13). Thus, the current DVI chips are inappropriate for the consumer market. In any case, it looks like Intel's proprietary algorithms have little chance against the industry support for MPEG, so Intel will have to play catch-up and implement MPEG (which it plans to do). Intel's current DVI chips do not have the horsepower to implement MPEG, and its next-generation chips aren't expected until late this year.

IIT, which was the first company to market a programmable video compression engine claimed to have MPEG capability (see µPR 10/30/91, p. 1), questions the viability of the CD-I market, noting that Nintendo computer-animation based games are usually better than the CD-I equivalents. IIT is focusing on communication-based markets and on MPEG authoring. IIT sees a chicken-and-egg problem with MPEG, and, like C-Cube, is working with PC companies to develop real-time compression chips.

Whether or not the 450 finds a huge market, C-Cube is clearly set up to be a force in the MPEG markets. By aggressively pursuing chip development in parallel with the completion of the compression standards and forging partnerships with major customers, C-Cube has taken a leading role. With time, multimedia will make video publishing as widespread as laser printers have made print publishing, and MPEG-II will make inroads into broadcast and cable TV. Unless it makes a significant error, C-Cube will have a piece of these lucrative pies. ♦