

## WP 2.5: A Video Decoder for H.261 Video Teleconferencing and MPEG Stored Interactive Video Applications

D. Brinthaupt, L. Letham, V. Maheshwari, J. Othmer, R. Spiwak, B. Edwards<sup>1</sup>, C. Terman<sup>1</sup>, N. Weste<sup>1</sup>

AT&T Bell Labs, Allentown, PA and Holmdel, NJ/TLW, Incorporated, Burlington, MA

This 450MOPS video decoder decompresses both H.261 and MPEG compressed video streams [1, 2]. The decoder accepts bit rates up to 4Mb/s and provides decoded frames of up to 352x288 pixels (CIF) at up to 30 frames/s operating at 45MHz.

The decoder places no restrictions on the H.261 bit streams. It decodes any combination of intra and predictive frames in QCIF or CIF format. In MPEG mode, it decodes any stream conforming to the MPEG "constrained parameters" (no more than 396 macroblocks per frame), including any combination of intra, predictive, and bidirectional frames with half-pixel motion vectors.

The architecture features a mix of dedicated hardware functions and programmable processors. Dedicated blocks provide performance (speed, power, and area) for those functions that are computationally expensive and well understood (e.g., color space conversion). Programmable processors with specialized architectures are used for functions that may need to be modified to suit different users or changes in the standards.

The decoder consists of 9 major blocks as shown in Figure 1. The host/serial interface block synchronizes the host bus and serial bus signals to the main input clock. It also contains the chip-reset logic and serves as the interface between the host and other decoder blocks. The 4kb compressed data FIFO is written into from either the host or serial bus.

The variable-length decoder (VLD) performs Huffman decoding on the bit stream using a micro-coded architecture for searching binary trees. The programs for H.261 and MPEG are stored in an on-chip 6kx32b ROM. A binary tree is traversed in response to the incoming bit stream until a leaf is encountered, whereupon a symbol representing the decoded bit sequence is transmitted through a FIFO to the symbol processor. Rather than simply returning to the root of the decoder tree, leaves are linked to other binary trees according to the hierarchical code tables specified in the standards. A data-bypass register allows embedded uncoded data to be passed directly to the symbol processor. Error recovery is handled in two ways. First, if an invalid node is entered, decoding is shifted to a special binary tree that searches for a predetermined bit pattern. Decoding proceeds only when this pattern is found. Second, the incoming bit stream is independently scanned for a start code. When the code is encountered, decoding is forced to the appropriate initialization point.

The symbol processor is a 45MIPS, 16b micro-programmed controller that interprets symbols from the VLD and issues commands to the memory controller and signal processor. The symbol processor is responsible for assembling run-length coded coefficients into DCT difference blocks, interpreting motion vectors and reconstructing the current frame from the decoded difference and previous frame information. It initiates and synchronizes the operation of all other modules. The symbol processor consists of two execution units: a traditional datapath with a 128-word register file, and a specialized datapath for assembling DCT blocks. This allows processing

DCT coefficients at a burst rate of 45M coefficients/s. The symbol processor executes 32b microcode from an on-chip 1.5k instruction RAM (IRAM). The microcode is downloaded into the IRAM through the host bus interface. The microcode also implements video stream status and control functions for the host program.

The single instruction multiple data (SIMD) signal processor performs the arithmetic functions, such as inverse discrete cosine transform (IDCT) and inverse quantization, on six 8x8 pixel blocks in parallel. Each of the six processing elements consists of a 152x9b input cache, a 144x16b register array, a 144x9b output cache, a 16b ALU, and a multiply/accumulate block. The six processors share an instruction execution unit, a 320x40b microcode ROM, and a 384x16b constant array for IDCT and inverse quantization coefficients. Each of the processors is capable of executing one instruction per clock cycle, for a total of 270MIPS at 45MHz. Originally designed for the companion video encoder, this block was adapted for use in the decoder by modifying the microcode ROM [3].

The memory controller provides a glue-free interface to external industry standard DRAMs and supports 60MB/s bandwidth. This hard-wired state machine performs four functions: write a macroblock into the DRAM from the signal processor; read a macroblock from the DRAM into the signal processor; read a fully decoded frame from the DRAM into the color converter while converting from 4:1:1 YCbCr to 4:2:2 YCbCr by vertical replication of chrominance data; and automatically perform CAS before RAS refresh cycles.

The color converter takes the 4:2:2 YCbCr data from the memory controller and optionally performs a programmable 3x3 matrix multiplication to generate 24b RGB values. The color converter has three multipliers so each pixel conversion requires three clock cycles. The output of the color converter, either YCbCr or RGB, passes through a lookup table (3x256x8b SRAM) where transformations such as gamma correction are performed, and then through a 32-pixel FIFO that is read from either the video output bus or the host.

The design methodology used for the decoder included extensive high-level modeling at two levels: a C++ behavioral model, and a set of clock-cycle-accurate C models at the block level. These models allowed microcode for both the symbol processor and the signal processor to be tested thoroughly enough so that the code worked in the first functional silicon, and were used to generate test vectors for gate-level simulations (and eventually for production test). Full-chip verification included running more than 5M vectors (1 vector per clock cycle) on a gate-level hardware accelerated simulator. Custom blocks were independently verified with full-timing simulations. The whole chip was also simulated in full-timing for several hundred clock cycles to verify the operation of major data paths.

The decoder layout is a hierarchical mix of full custom and standard cell blocks. Chip characteristics are given in Table 1. A die photograph is shown in Figure 2. Silicon is fully functional at 45MHz and 5.0V.

References

- [1] "Video Codec for Audio Visual Services at px64 kb/s", CCITT Recommendation H.261, 1990.
- [2] "Coding of Moving Pictures and Associated Audio for Digital Storage Media up to about 1.5 Mb/s", Committee Draft of Standard DIS 11172 ISO/MPEG 90/176, Dec. 1990.
- [3] Rao, S., et al., "A Real-Time P\*64/MPEG Video Encoder Chip", ISSCC DIGEST OF TECHNICAL PAPERS, pp. 32-33, Feb. 1993.

Die Size	12.0x12.0mm <sup>2</sup>
Technology	0.75μm 2-level metal CMOS
Number of devices	1.2M
Clock Rate	45MHz
Power consumption	2.5W @ 5V, 45MHz
Computational power	450MOPS
Video rate	Up to 30 frames/s
Image size	Up to CIF (352x288) motion Up to 1024x1024 still frame
Input bit rate	Up to 4Mb/s

Table 1: Video decoder characteristics.




Figure 1: Video decoder block diagram.

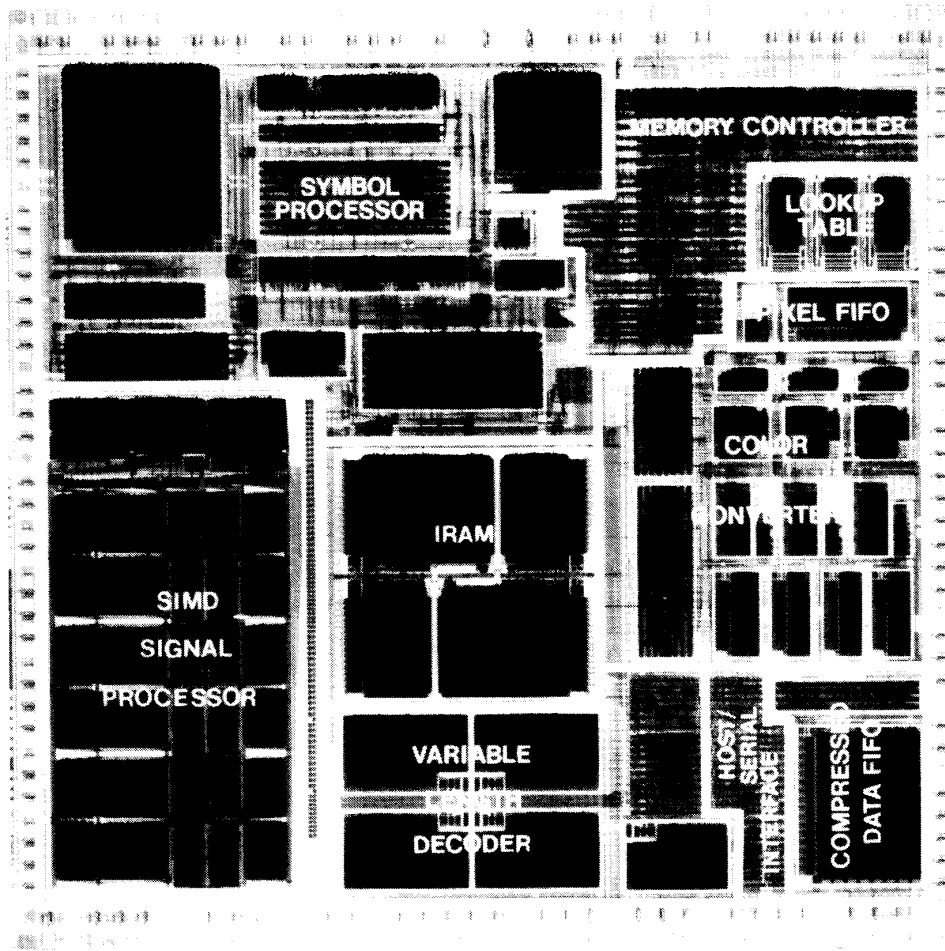


Figure 2: Video decoder die micrograph.