
The MVP: A Highly-Integrated Video Compression Chip

Robert J. Gove

Texas Instruments, Inc.
Dallas, Texas 75265

ABSTRACT

We introduce a new highly-integrated processing chip for performing a variety of functions, however this chip is particularly well suited for video compression algorithms. Applications include multimedia PCs, virtual reality 3D graphics, full-duplex videoconferencing, HDTV, and color hardcopy. We have architected the Multimedia Video Processor, or MVP, to provide a yet unattainable level of *performance* from a single chip, although with the *programmability* typically found in today's general-purpose computers. While advanced semiconductor design and process techniques have been used for its design, the *key* to the advantage of this component lies in optimization of the architecture for real-time video and graphics processing. This paper will analyze video compression application requirements, describe the MVP architecture, and pose its potential as a very capable solution for a wide range of markets.

INTRODUCTION

The computer and consumer video industries are pursuing varied paths to offer cost-effective computing products which provide new forms of information and entertainment. Products are emerging from cable TV delivery of interactive digital movies to digital mobile offices. Digital compression and video processing at a reasonable cost are spurring this revolution. While algorithm developments have been important, most of the enabling advances lie in the availability of high-density memory and high-performance processing ICs. With the pending general availability of the Multimedia Video Processor, or MVP, in 1994, a yet unattained level of digital signal processing performance will be available and with all the flexibility of present day programmable computers. Standard-based videoconferencing and playback of compressed digital video and audio (using Px64, JPEG or MPEG "multi-standard" codecs systems) with a single MVP processor will be possible, as well as codecs with yet-to-be-defined algorithms like model-based compression. However, not only will the MVP support compression, it will also handle processing of high-resolution video, full-motion video processing from sources like camcorders, digital audio processing, hardcopy raster image processing, and 3D graphics, and all under software control and generation. From this wide range of functions, we calculated that several billion operations per second are required to provide video-based applications on the desktop. Current and soon to appear desktop host processors like X86, Pentium, Alpha, and MIPS do not have the computational power to meet these demands.

KEYS TO THE MVP ARCHITECTURE

The MVP's unique architecture and computational power enables users to integrate these varied functions on a single processing component. The keys to obtaining both exceptional processing speeds and fully-programmable features with the MVP include the use of:

- (1) *an efficient parallel processing architecture,*
- (2) *fast pixel processing tuned to image, video, and graphics processing,*
- (3) *intelligent control of image data flow throughout the architecture,*
- (4) *single-chip integration without slower chip-to-chip communications.*




Figure 1: MVP Block Diagram:
(A Single-Chip Parallel Processor)

ALGORITHM-DIRECTED ARCHITECTURE DEFINITION

Processing Requirements

Today's proposed international video compression standards use common frequency domain, quantization, and entropy coding techniques to (de)compress small portions (8x8) of each image. While these functions demand a great deal from the encoder/decoder, many other varied functions remain, each with dynamic requirements which vary based on the type of image compressed as well as the channel rate required to maintain real-time operation. For optimal efficiency a processor must adapt to these dynamic needs. A typical average of the processing demands of the Px64 video-conferencing standard appears in the following table.

RISC vs. MVP-PP Processing Requirements for Px64

Px64 (H.261) FULL-DUPLEX, FULL-CIF, 30Hz Functions	RISC Execution Speed (average % of time) *	MVP Execution Speed (average % of time)	Speed-up of MVP-PP vs. RISC Processors
Motion Estimation - Block Matching (encode)	0.51	0.29	14
Encoding Decisions - (1) Inter w/motion vectors, (2) Inter w/coded diff., (3) Intra	0.034	0.039	7
Loop Filtering (both)	0.092	0.116	6
Difference image (current - predicted)	0.18	0.013	9
Fast DCT (encode)	0.062	0.077	6
Threshold/Quantization/Zig-Zag Run-length	0.042	0.071	5
Bitstream Encode	0.014	0.045	3
IDCT (both)	0.161**	0.226	6
Reconstruction (both) (predicted + diff. image)	0.052	0.077	5
Bitstream Decode & Dequantization (decode)	0.018	0.045	3
TOTAL CYCLES (MIPS)	1.00 = 1,193 MIPS	1.00 = 155 PP MIPS ***	AVERAGE SPEED-UP = 7.7

* Multiply counted a one instruction even though most RISCs require many cycles.

** If the "Truncated-IDCT" algorithm was used, IDCTs speed-up again (see later).

*** The total is equivalent to 3 MVP-PP processors (see below PP section).

**** Audio standards concurrently execute on the MVP-MP (see below MP section).

As we studied the computational requirements for motion estimation (51%) and DCTs (22%) it became quite apparent that a programmable image processor must excel at these functions. It is important to recognize that what's done poorly in a processor can dominate its performance. Since most architectural improvements would not uniformly accelerate all functions uniformly, we looked for special architectural features for these critical functions, while maintaining enough flexibility to benefit a larger class of algorithms. In final analysis, a much more uniform distribution of computational loading resulted after the changes.

As seen in the table, the programmable image processor must perform many other functions well, including: bit manipulation and table look ups for entropy encoding, and

multiply and accumulate for various types of filtering operations. To obtain good image quality at any channel rate and 30 frames per second, the image processor must compute over 1.2 billion operations per second (BOPS).

The addition of audio compression (which requires higher precision integer and possibly floating point algorithms) and network communication, necessary for video conferencing (G.728 or G.711, H.242, H.230, H.221), further increases the scope of computational requirements. Reducing the system cost, we propose to include support in the architecture for the required non-standard functions like color space conversion (YCrCb to RGB), decimation of the source image to CIF resolution and variable scaling of the decompressed sequence. Complete implementation of compression applications such as video-conferencing requires over 2 BOPS of the programmable image processor.

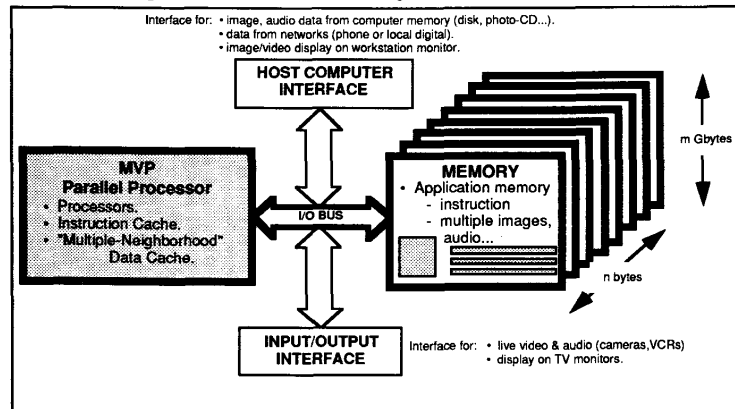
ARCHITECTURE CHOICES

We considered several candidate parallel architectures for implementation of this single-chip video processor [Gove-92, Guttag-92]. An architecture with a mix of dedicated and programmable processors was initially evaluated, then subsequently discounted when no single dominant function was found that was necessary almost all of the time. Besides, we predicted that by the time the chip was completed, that a new *important* algorithm would emerge. From the standpoint of loss of silicon efficiency by dedicated resources to any one function (like a DCT), we felt compelled to seek a general-purpose well-balanced system solution. Several other candidates existed, however the mix of algorithms and practical implementation limitations focused us on SIMD and MIMD architectures. These differ by the autonomy of the processors functions with MIMD -- a desirable feature for any data dependent algorithm operating in parallel.

With MIMD desirable, the choice of a processor and memory interconnection architecture remained. Pipelined, shared bus memory, communication port (mesh/array/hypercube), and crossbar fully-shared memory were considered. Pipeline memory and processors (systolic arrays) are typically used for video, however they're too restrictive in the sense that one must a priori know the size of the memory and dynamics of the algorithm to prevent data contention and processor stalls. With our varied needs, this would lead to inefficiencies. A shared-bus memory structure would also have bottleneck problems with highly variable instruction and data streams and moving of results from one processor to the other. The n-way connected communication port requires a very ordered flow of data, like a systolic or wavefront flow of data, or the application of a pixel per processor (not practical in a single chip). This approach works for large arrays of simple processors which can operate uniformly on images, however we wanted more complex processors which could adapt to varying types of data, from bit graphics to floating-point representations. The crossbar fully-shared memory is ideally suited to these needs, minimizing contention, data movement and providing flexibility for many types of algorithms. In fact, since the crossbar operations at the processor instruction rates, this architecture can functionally emulate the other approaches (pipeline, shared bus...).

We not only wanted to provide this order of magnitude performance increase, but the goal was to apply a *traditional* computer model of programmable processing and a large memory to applications with integrated image, graphics, video and audio processing, or image computing. As shown in Figure #2 titled "*MVP System Architecture*", replacing the processing and memory pipeline of conventional video systems with the single video processor and large memory system model yields tremendous application flexibility. In effect the system can re-configure itself with software from video conferencing to playing CD movies, just as a PC would re-configure from a spreadsheet to a video game.

Figure 2: The MVP "System" Architecture.



THE MVP ARCHITECTURE

The Multimedia Video Processor, or MVP, represents the next-generation of digital signal processors. The MVP can be technically described as a single-chip crossbar shared memory heterogeneous MIMD multiprocessor. It combines RISC and advanced DSP processing in one parallel architecture with unique features for each. Current RISC processors typically use instruction pipelining, numerous registers and a detached floating point processor. On the other hand, current DSPs are optimized for one dimensional multiply-accumulate functions. Newer DSPs have floating-point capabilities, yet most imaging and video only needs integer operations. DSPs usually have fewer registers than RISC and have direct memory accesses (DMA) with limited capabilities.

The MVP combines the best features of RISC and DSP in parallel and adds other features to offer unprecedented *Power* and *Flexibility*. The heart of an image or video chip is its capability to process 2D signals. The MVP has features for 2D DSP-like processing, including multiply-accumulate operations. The on-chip memory and register characteristics of the MVP were optimized for image computing algorithms, preventing time consuming cache misses or swapping of register contents. Multidimensional external memory access and double buffering minimizes the typical memory bottleneck of current DSP solutions. An internal memory crossbar provides extremely efficient synchronization and communication of multiple processors. A very high-performance RISC processor is integrated on the chip, providing intelligent control of the DSP-like processors. Also integrated into the chip, a new floating-point architecture can act as a co-processor to any of the DSP-like processors or the RISC processor. By analysis of the algorithms, the required mix of integer ops to floating-point ops was somewhere between 8:1 and 4:1 -- a balance which the MVP supports. The entire collection of processors and memory is configured as a MIMD architecture for ease of programming and high performance for all image and video computing applications. This MIMD data and control supports both data

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.