


HEWLETT-PACKARD JOURNAL

April 1995



© Copr. 1949-1998 Hewlett-Packard Co.

 **HEWLETT®
PACKARD**

AMAZON-1021
8,204,134

Articles

- 6 **A Low-Cost, High-Performance PA-RISC Workstation with Built-In Graphics, Multimedia, and Networking Capabilities**, *by Roger A. Pearson*
- 12 **The PA 7100LC Microprocessor: A Case Study of IC Design Decisions in a Competitive Environment**, *by Mick Bass, Patrick Knebel, David W. Quint, and William L. Walker*
- 23 **Design Methodologies for the PA 7100LC Microprocessor**, *by Mick Bass, Terry W. Blanchard, D. Douglas Josephson, Duncan Weir, and Daniel L. Halperin*
- 36 **An I/O System on a Chip**, *by Thomas V. Spencer, Frank J. Lettang, Curtis R. McAllister, Anthony L. Riccio, Joseph F. Orth, and Brian K. Arnold*
- 43 **An Integrated Graphics Accelerator for a Low-Cost Multimedia Workstation**, *by Paul Martin*
- 51 **HP Color Recovery Technology**, *by Anthony C. Barkans*
- 52 **True Color**
- 60 **Real-Time Software MPEG Video Decoder on Multimedia-Enhanced PA 7100LC Processors**, *by Ruby B. Lee, John P. Beck, Joel Lamb, and Kenneth E. Severson*
- 66 **Overview of the Implementation of the PA 7100LC Multimedia Enhancements**
- 69 **HP TeleShare: Integrating Telephone Capabilities on a Computer Workstation**, *by S. Paul Tucker*
- 72 **Caller-ID**
- 73 **Call Progress, DTMF Tones, and Tone Detection**
- 75 **Product Design of the Model 712 Workstation and External Peripherals**, *by Arlen L. Roesner*

Editor, Richard P. Dolan • Associate Editor, Charles L. Leath • Publication Production Manager, Susan E. Wright • Illustration, Renée D. Pighini

Advisory Board: Rajeev Badyal, *Integrated Circuit Business Division, Fort Collins, Colorado* • Thomas Beecher, *Open Systems Software Division, Chelmsford, Massachusetts* • Steven Brittenham, *Disk Memory Division, Boise, Idaho* • William W. Brown, *Integrated Circuit Business Division, Santa Clara, California* • Harry Chou, *Microwave Technology Division, Santa Rosa, California* • Rajesh Desai, *Commercial Systems Division, Cupertino, California* • Kevin G. Ewert, *Integrated Systems Division, Sunnyvale, California* • Bernhard Fischer, *Böblingen Medical Division, Böblingen, Germany* • Douglas Gennetten, *Greeley Hardcopy Division, Greeley, Colorado* • Gary Gordon, *HP Laboratories, Palo Alto, California* • Matt J. Harline, *Systems Technology Division, Roseville, California* • Kiyoyasu Hiwada, *Hachioji Semiconductor Test Division, Tokyo, Japan* • Bryan Hoog, *Lake Stevens Instrument Division, Everett, Washington* • Roger L. Jungerman, *Microwave Technology Division, Santa Rosa, California* • Ruby B. Lee, *Networked Systems Group, Cupertino, California* • Alfred Maute, *Waldbronn Analytical Division, Waldbronn, Germany* • Dona L. Miller, *Worldwide Customer Support Division, Mountain View, California* • Michael P. Moore, *VXI Systems Division, Loveland, Colorado* • Shelley I. Moore, *San Diego Printer Division, San Diego, California* • M. Shahid Mujtaba, *HP Laboratories, Palo Alto, California* • Steven J. Narciso, *VXI Systems Division, Loveland, Colorado* • Danny J. Oldfield, *Colorado Springs Division, Colorado Springs, Colorado* • Garry Orsolini, *Software Technology Division, Roseville, California* • Han Tian Phua, *Asia Peripherals Division, Singapore* • Ken Poulton, *HP Laboratories, Palo Alto, California* • Günter Riebesell, *Böblingen Instruments Division, Böblingen, Germany* • Marc Sabatella, *Software Engineering Systems Division, Fort Collins, Colorado* • Michael B. Saunders, *Integrated Circuit Business Division, Corvallis, Oregon* • Philip Stenton, *HP Laboratories, Bristol, Bristol, England* • Stephen R. Undy, *Systems Technology Division, Fort Collins, Colorado* • Jim Willits, *Network and System Management Division, Fort Collins, Colorado* • Koichi Yanagawa, *Kobe Instrument Division, Kobe, Japan* • Dennis C. York, *Corvallis Division, Corvallis, Oregon* • Barbara Zimmer, *Corporate Engineering, Palo Alto, California*

-
-
- 79 **Development of a Low-Cost, High-Performance, Multiuser Business Server System**, by *Dennis A. Bowers, Gerard M. Enkerlin, and Karen L. Murillo*
-
- 85 **HP Distributed Smalltalk: A Tool for Developing Distributed Applications** by *Eileen Keremitsis and Ian J. Fuller*
-
- 86 **Object Management Group**
-
- 93 **A Software Solution Broker for Technical Consultants**, by *Manny Yousefi, Adel Ghoneimy, and Wulf Rehder*
-
- 98 **HP Software Solution Broker Accessible Products**
-
- 102 **Bugs in Black and White: Imaging IC Logic Levels with Voltage Contrast**, by *Jack D. Benzel*
-
- 107 **Component and System Level Design-for-Testability Features Implemented in a Family of Workstation Products**, by *Bulent I. Dervisoglu and Michael Ricchetti*
-

Departments

- 4 **In this Issue**
5 **Cover**
5 **What's Ahead**
114 **Authors**

The Hewlett-Packard Journal is published bimonthly by the Hewlett-Packard Company to recognize technical contributions made by Hewlett-Packard (HP) personnel. While the information found in this publication is believed to be accurate, the Hewlett-Packard Company disclaims all warranties of merchantability and fitness for a particular purpose and all obligations and liabilities for damages, including but not limited to indirect, special, or consequential damages, attorney's and expert's fees, and court costs, arising out of or in connection with this publication.

Subscriptions: The Hewlett-Packard Journal is distributed free of charge to HP research, design and manufacturing engineering personnel, as well as to qualified non-HP individuals, libraries, and educational institutions. Please address subscription or change of address requests on printed letterhead (or include a business card) to the HP headquarters office in your country or to the HP address on the back cover. When submitting a change of address, please include your zip or postal code and a copy of your old label. Free subscriptions may not be available in all countries.

The Hewlett-Packard Journal is available online via the World-Wide Web (WWW) and can be viewed and printed with Mosaic. The uniform resource locator (URL) for the Hewlett-Packard Journal is <http://www.hp.com/hpj/Journal.html>.

Submissions: Although articles in the Hewlett-Packard Journal are primarily authored by HP employees, articles from non-HP authors dealing with HP-related research or solutions to technical problems made possible by using HP equipment are also considered for publication. Please contact the Editor before submitting such articles. Also, the Hewlett-Packard Journal encourages technical discussions of the topics presented in recent articles and may publish letters expected to be of interest to readers. Letters should be brief, and are subject to editing by HP.

Copyright © 1995 Hewlett-Packard Company. All rights reserved. Permission to copy without fee all or part of this publication is hereby granted provided that 1) the copies are not made, used, displayed, or distributed for commercial advantage; 2) the Hewlett-Packard Company copyright notice and the title of the publication and date appear on the copies; and 3) a notice appears stating that the copying is by permission of the Hewlett-Packard Company.

Please address inquiries, submissions, and requests to: Editor, Hewlett-Packard Journal, 3000 Hanover Street, Palo Alto, CA 94304 U.S.A.

Real-Time Software MPEG Video Decoder on Multimedia-Enhanced PA 7100LC Processors

With a combination of software and hardware optimizations, including the availability of PA-RISC multimedia instructions, a software video player running on a low-end workstation is able to play MPEG compressed video at 30 frames/s.

by Ruby B. Lee, John P. Beck, Joel Lamb, and Kenneth E. Severson

Traditionally, computers have improved productivity by helping people compute faster and more accurately. Today, computers can further improve productivity by helping people communicate better and more naturally. Towards this end, at Hewlett-Packard we have looked for more natural ways to integrate communication power into our desktop machines, which would allow a user to access distributed information more easily and communicate with other users more readily.

We felt that adding audio, images, and video information would enrich the information media of text and graphics normally available on desktop computers such as workstations and personal computers. However, for such enriched multimedia communications to be useful, it must be fully integrated into the user's normal working environment. Hence, as the technology matured we decided to integrate increasing levels of multimedia support into both the user interface and the basic hardware platform.

In terms of user interface, we integrated a panel of multimedia icons into the HP VUE standard graphical user interface, which comes with all HP workstations. These multimedia icons are part of the HP MPower product.¹ HP MPower enables a workstation user to receive and send faxes, share printers, access and manipulate images, hear and send voice and CD-quality stereo audio, send and receive multimedia email, share an X window or an electronic whiteboard with other distributed users, and capture and play back video sequences. The HP MPower software is based on a client/server model, in which one server can service around 20 clients, which can be workstations or X terminals.

In terms of hardware platforms, we integrated successive levels of multimedia support into the baseline PA-RISC workstations.^{2,3,4} First, we integrated support for all the popular image formats such as JPEG (Joint Photographic Experts Group)[†] compressed images.⁵ Then, we added hardware and software support for audio, starting with 8-kHz voice-quality audio, followed by support for numerous audio formats including A-law, μ -law, and 16-bit linear mode, with up to 48-kHz mono and stereo. This allowed high-fidelity,

44.1-kHz stereo, 16-bit CD-quality audio to be recorded, manipulated, and played back on HP workstations. At the same time, we supported uncompressed video capture and playback.

In January 1994, HP introduced HP MPower 2.0 and the entry-level enterprise workstation, the HP 9000 Model 712, which is based on the multimedia-enhanced PA-RISC processor known as the PA 7100LC.^{6,7,8} The video player integrated in the MPower 2.0 product is the first product that achieves real-time MPEG-1 (Moving Picture Experts Group)⁹ video decompression via software running on a general-purpose processor. Typically real-time MPEG-1 decompression is achieved via special-purpose chips or boards. Previous attempts at software MPEG-1 decompression did not attain real-time rates.¹⁰ The fact that this is achieved by the low-end Model 712 workstation is significant.

In this paper, we discuss the support of MPEG-compressed video as a new (video) data type. In particular, we discuss the technology that enables the video player integrated into the HP MPower 2.0 product to play back MPEG-compressed video at real-time rates of up to 30 frames per second.

Digital Video Standards

We decided to focus on the MPEG digital video format because it is an ISO (International Standards Organization) standard, and it gives the highest video fidelity at a given compression ratio of any of the formats that we evaluated. MPEG also has broad support from the consumer electronics, telecommunications, cable, and computer industries. The high compression capability of MPEG translates into lower storage costs and less bandwidth needed for transmitting video on the network. These characteristics make MPEG an ideal format for addressing the need for detail in the video used in technical workstation markets and computer-based training in commercial workstation markets.

MPEG is one of several algorithmically related standards shown in Fig. 1. All of these digital video compression standards use the discrete cosine transform (DCT) as a fundamental component of the algorithm. Alternatives to discrete cosine-based algorithms that we looked at include vector quantization, fractals, and wavelets. Vector quantization

[†] JPEG is an international digital image compression standard for continuous-tone (multi-level) still images (grayscale and color).




Fig. 1. Digital video standards based on the discrete cosine transform.

algorithms are popular on older computer architectures because they require less computing power to decompress, but this advantage is offset by poorer image quality at low bandwidth (high compression) compared to MPEG for practical vector quantization methods. Algorithms based on wavelet and fractal technology have the potential to deliver video fidelity comparable to MPEG, but there is presently a lack of industry consensus on standardization, a key requirement for our use.

Another advantage of a high-performance implementation of MPEG is the ability to leverage the improvements to the other DCT-based algorithms. Although the relationships shown in Fig. 1 do not represent a true hierarchy of algorithms is useful for illustrating increased complexity as one moves from JPEG to MPEG-2, or from H.261 to MPEG-2.

All of these formats have much in common, such as the use of the DCT for encoding. The visual fidelity of the algorithms was the key selection criterion and not ease of implementation or performance on existing hardware.

Although JPEG supports both lossy and lossless compression, the term JPEG is typically associated with the lossy specification.[†] The primary goal of JPEG is to achieve high compression of photographic images with little perceived loss of image fidelity. Although it is not an ISO standard, by convention, a sequence of JPEG lossy images to create a digital video sequence is called motion JPEG, or MJPEG.

H.261 is a digital video standard from the telecommunications standards body ITU-TSS (formerly known as CCITT). H.261 is one of a suite of conferencing standards that make up the umbrella H.320 specification. H.261 is often referred to as P*64 (where P is an integer) because it was designed to fit into multiples of 64 kbits/s bandwidth. The first frame

[†] In lossless compression, decompressed data is identical to the original image data. In lossy compression, decompressed data is a good approximation of the original image data.

(image) of an H.261 sequence is for all practical purposes a highly compressed lossy JPEG image. Subsequent frames are built from image fragments (blocks) that are either JPEG-like or are differences from the image fragments in previous frames. Most video sequences have high frame-to-frame coherence. This is especially true for video conferencing. Because the encoding of the movement of a piece of an image requires less data than an equivalent JPEG fragment, H.261 achieves higher visual fidelity for a given bandwidth than does motion JPEG. Since the encoding of the differences is always based on the previous frames, the technique is called *forward differencing*.

The MPEG-1 specification goes even further than H.261 in allowing sophisticated techniques to achieve high fidelity with fewer bits. In addition to forward differencing, MPEG-1 allows backward differencing (which relies on information in a future frame) and averaging of image fragments. (Forward and backward differencing are described in more detail in the next section.) MPEG-1 achieves quality comparable to a professionally reproduced VHS videotape even at a single-speed CD-ROM data rate (1.5 Mbits/s).^{9,11} MPEG-1 also specifies encodings for high-fidelity audio synchronized with the video.

MPEG-2 contains additional specifications and is a superset of MPEG-1. The new features in MPEG-2 are targeted at broadcast television requirements, such as support for frame interleaving similar to analog broadcast techniques. With widespread deployment of MPEG-2, the digital revolution for video may be comparable to the digital audio revolution of the last decade.

The approximate bandwidths required to achieve a level of subjective visual fidelity for motion JPEG, H.261, MPEG-1, and MPEG-2 are shown in Fig. 2. Motion JPEG will primarily be used for cases in which accurate frame editing is important such as video editing. H.261 will be used primarily for video conferencing, but it also has potential for use in video mail. MPEG-1 and MPEG-2 will be used for publishing, where fidelity expectations have been set by consumer analog video tapes, computer-based training, games, movies on CD, and video on demand.

MPEG Compression

MPEG has two classes of frames: intracoded and non-intracoded frames (see Fig. 3). Intracoded frames, also called *I-frames*, are compressed by reducing spatial redundancy within the frame itself. I-frames do not depend on comparisons with past or "reference" frames. They use JPEG-type compression for still images.⁵




Fig. 2. Compressed video bandwidth versus subjective visual fidelity. The ideal format achieves exceptional visual fidelity at the lowest bandwidth.




Fig. 3. MPEG frame sequencing.

Nonintra-coded frames are further divided into *P-frames* and *B-frames*. *P-frames* are predicted frames based on comparisons with an earlier reference frame (an intra-coded or predicted frame). By considering temporal redundancy in addition to spatial redundancy, *P-frames* can be encoded with fewer bits. *B-frames* are bidirectionally predicted frames that require one backward reference frame and one forward reference frame for prediction. A reference frame can be an *I-frame* or a *P-frame*, but not a *B-frame*. By detecting the motion of blocks from both a frame that occurred earlier and a frame that will be played back later in the video sequence, *B-frames* can be encoded in fewer bits than *I-* or *P-frames*.

Each frame is divided into macroblocks of 16 by 16 pixels for the purposes of motion estimation† in MPEG compression and motion compensation in MPEG decompression. A frame with only *I-blocks* is an *I-frame*, whereas a *P-frame* has *P-blocks* or *I-blocks*, and a *B-frame* has *B-blocks*, *P-blocks*, or *I-blocks*. For each *P-block* in the current frame, the block in the reference frame that matches it best is identified by a motion vector. Then the differences between the pixel values in the matching block in the reference frame and the current block in the current frame are encoded by a discrete cosine transform.

The color space used is the YCbCr color representation rather than the RGB color space, where *Y* represents the luminance (or brightness) component, and *Cb* and *Cr* represent the chrominance (or color) components. Because human perception is more sensitive to luminance than to chrominance, the *Cb* and *Cr* components can be subsampled in both the *x* and *y* dimensions. This means that there is one *Cb* value and one *Cr* value for every four *Y* values. Hence, a 16-by-16 macroblock contains four 8-by-8 blocks of *Y*, and only one 8-by-8 block of *Cb* and one 8-by-8 block of *Cr* values (see Fig. 4). This is a reduction from the twelve 8-by-8 blocks (four for each of the three color components) if *Cb*

and *Cr* were not subsampled. The six 8-by-8 blocks in each 16-by-16 macroblock then undergo transform coding.

Transform coding concentrates energy in the lower frequencies. The transformed data values are then quantized by dividing by the corresponding quantization coefficient. This results in discarding some of the high-frequency values, or lower-frequency but low-energy values, since these become zeros. Both transform coding and quantization enable further compression by run-length encoding of zero values.

Finally, the nonzero coefficients of an 8-by-8 block used in the discrete cosine transform can be encoded via variable-length entropy encoding such as Huffman coding. Entropy encoding basically removes coding redundancy by assigning the code words with the fewest number of bits to those coefficients that occur most frequently.




Fig. 4. Subsampling of the chrominance components (*Cb*, *Cr*) with respect to the luminance (*Y*) component.

† Motion estimation uses temporal redundancy to estimate the movement of a block from one frame to the next.

MPEG Decompression

MPEG decompression reverses the functional steps taken for MPEG compression. There are six basic steps involved in MPEG decompression.

1. The MPEG header is decoded. This gives information such as picture rate, bit rate, and image size.
2. The video data stream is Huffman or entropy decoded from variable-length codes into fixed-length numbers. This step includes run-length decoding of zeros.
3. Inverse quantization is performed on the numbers to restore them to their original range.
4. An inverse discrete cosine transform is performed on the 8-by-8 blocks in each frame. This converts from the frequency domain back to the original spatial domain. This gives the actual pixel values for I-blocks, but only the differences for each pixel for P-blocks and B-blocks.
5. Motion compensation is performed for P-blocks and B-blocks. The differences calculated in step 4 are added to the pixels in the reference block as determined by the motion vector for P-blocks and to the average of the forward and backward reference blocks for B-blocks.
6. The picture is displayed by doing a color conversion from YCbCr coordinates to RGB color coordinates and writing to the frame buffer.

Methodology

Our philosophy was to improve the algorithms and tune the software first, resorting to hardware support only if necessary. We set a goal of 10 to 15 frames/s for software MPEG video decompression because this is the rate at which motion appears smooth rather than jerky.

We started by measuring the performance of the MPEG software we had purchased. This software initially took two seconds to decode one frame (0.5 frame/s) on an older 50-MHz Model 720 workstation. This decoding was for video only and did not include audio. Profiling indicated that the inverse discrete cosine transform (step 4) took the largest chunk of the execution time, followed by display (step 6), followed by motion compensation (step 5). The decoding of the MPEG headers was insignificant.

With this data we set out to optimize every step in the MPEG decompression software. After we applied all the algorithm enhancements and software tuning, we measured the MPEG decode software again. While we had achieved an order of magnitude improvement, the rate of 4 to 5 frames/s was not sufficient to meet our goal.

Hence, we looked at possible multimedia enhancements to the basic PA-RISC processor and other system-level enhancements that would not only speed up MPEG decoding, but also be generally useful for improving performance in other computations. In addition, any chip enhancements we added could not adversely impact the design schedule, complexity, cycle time, and chip size of the PA-RISC processor we were targeting, the PA 7100LC, which was already deep into its implementation phase at the time. The PA 7100LC is described in detail in the article on page 12.

We approached this problem by studying the distribution of operations executed by the software MPEG decoder. Then,

we found ways to reduce the execution time of the most frequent operation sequences. The application of algorithm enhancements, software tuning, and projected hardware enhancements was iterated until we attained our goal of being able to decompress at a rate greater than 15 frames/s via software.

Algorithm and Software Optimizations

In terms of MPEG video algorithms, we improved on the Huffman decoder, the motion compensation, and the inverse discrete cosine transform. A faster Huffman decoder based on a hybrid of table lookup and tree-based decoding is used. The lookup table sizes were chosen to reduce cache misses. For motion compensation, we sped up the pixel averaging operations.

For the inverse discrete cosine transform, we use a faster Fourier transform, which significantly reduces the number of multiplies for each two-dimensional 8-by-8 inverse discrete cosine transform. In addition, we use the fact that the 8-by-8 inverse transform matrices are frequently sparse to further reduce the multiplies and other operations required.

The MPEG audio decompression is also done in software. This algorithm was improved by using a 32-point discrete cosine transform to speed up the subband filtering.¹²

In terms of software tuning, we "flattened" the code to reduce the number of procedure calls and returns, and the frequent building up and tearing down of contexts present in the original MPEG code. We also did "strength reductions" like reducing multiplications to simpler operations such as shift and add or table lookup.

The last column of Table I shows the percentage of execution time spent in each of the six MPEG decompression steps after the algorithm and software tuning improvements were made. The first two columns of Table I show the millions of instructions executed in each of the six decompression steps and the percent of the total instructions executed (path length) each step represents. The input video sequence was an MPEG-compressed clip of a football game. The total time taken was 7.45 seconds on an HP 9000 Model 735 99-MHz PA-RISC workstation, with 256K bytes of instruction cache and 256K bytes of data cache.

Table I
Instructions and Time Spent in each MPEG Decompression Step on an HP 9000 Model 735

	Millions of Instructions	Path Length (%)	Time (%)
Header decode	0.6	0.1	0.1
Huffman decode	55.3	10.2	7.5
Inverse quantization	8.7	1.6	2.4
Inverse discrete cosine transform	206.5	38.3	38.7
Motion compensation	79.9	14.8	18.3
Display	188.7	35.0	33.0
Total	539.7	100.0	100.0

The largest slice of execution time (38.7%) and the largest chunk of instructions executed (38.3%) were still the inverse discrete cosine transform. We studied the frequencies of generic operations in this group and attempted to execute them faster. This resulted in new PA-RISC processor instructions for accelerating multimedia software.

PA-RISC Processor Enhancements

The new processor multimedia instructions implemented in the PA 7100LC processor allow simple arithmetic operations to be executed in parallel on subword data in the standard integer data path. In particular, the integer ALU is partitioned so that it can execute a pair of arithmetic operations in a single cycle with a single instruction. The arithmetic operations accelerated in this way are add, subtract, average, shift left and add, and shift right and add. The latter two operations are effective in implementing multiplication by constants.

PA-RISC Multimedia Extensions 1.0. The PA 7100LC PA-RISC processor chip contains some instructions that operate independently and in parallel on two 16-bit data fields within a 32-bit register. These operations are independent in that bits carried or shifted out of one of the fields never affects the result in the other field. These operations occur in parallel in that a single instruction computes both 16-bit fields of the result. Table II summarizes these instructions.

HADD does two parallel 16-bit additions on the left and the right halves of registers *ra* and *rb*, placing the two 16-bit results into the left and right halves of register *rt*.

HSUB does two parallel 16-bit subtractions on the left and right halves of registers *ra* and *rb*, placing the two 16-bit results into the left and right half of register *rt*.

Both HADD and HSUB perform modulo arithmetic (modulus 2^{16}), that is, the result wraps around from the largest number back to the smallest number and vice versa. This is the usual mode of operation of twos complement adders when overflow is ignored.

HADD and HSUB also have two saturation arithmetic options. With the signed saturation option, HADD.ss, both operands and the result are considered signed 16-bit integers. If the result cannot be represented as a signed 16-bit integer, it is clipped to the largest positive value ($2^{15}-1$) if positive overflow occurs, or it is clipped to the smallest negative value (-2^{15}) if negative overflow occurs.

With the unsigned saturation option, HADD.us, the first operand (*ra*) is considered an unsigned 16-bit integer, the second operand (*rb*) is considered a signed 16-bit integer, and the result (in *rt*) is considered an unsigned 16-bit integer. If the result cannot be represented as an unsigned 16-bit integer, it is clipped to the largest unsigned value ($2^{16}-1$) if positive overflow occurs, or it is clipped to the smallest unsigned value (0) if negative overflow occurs.

The signed saturation and unsigned saturation options for parallel halfword subtraction are defined similarly.

HAVE, or halfword average, gives the average of each pair of halfwords in *ra* and *rb*. It takes the sum of parallel halfwords and does a right shift of one bit before storing each 16-bit result into *rt*. During the one-bit right shift, the carry is

Table II
PA-RISC Multimedia Instructions in PA 7100LC
ra contains *a1*; *a2*
rb contains *b1*; *b2*
rt contains *t1*; *t2*

Instruction	Parallel Operation
HADD <i>ra,rb,rt</i>	$t1 = (a1+b1) \bmod 2^{16};$ $t2 = (a2+b2) \bmod 2^{16};$
HADD.ss <i>ra,rb,rt</i>	$t1 = \text{IF } (a1+b1) > (2^{15}-1) \text{ THEN } (2^{15}-1)$ $\text{ELSEIF } (a1+b1) < -2^{15} \text{ THEN } (-2^{15})$ $\text{ELSE } (a1+b1);$ $t2 = \text{IF } (a2+b2) > (2^{15}-1) \text{ THEN } (2^{15}-1)$ $\text{ELSEIF } (a2+b2) < -2^{15} \text{ THEN } (-2^{15})$ $\text{ELSE } (a2+b2);$
HADD.us <i>ra,rb,rt</i>	$t1 = \text{IF } (a1+b1) > (2^{16}-1) \text{ THEN } (2^{16}-1)$ $\text{ELSEIF } (a1+b1) < 0 \text{ THEN } 0$ $\text{ELSE } (a1+b1);$ $t2 = \text{IF } (a2+b2) > (2^{16}-1) \text{ THEN } (2^{16}-1)$ $\text{ELSEIF } (a2+b2) < 0 \text{ THEN } 0$ $\text{ELSE } (a2+b2);$
HSUB <i>ra,rb,rt</i>	$t1 = (a1-b1) \bmod 2^{16};$ $t2 = (a2-b2) \bmod 2^{16};$
HSUB.ss <i>ra,rb,rt</i>	$t1 = \text{IF } (a1-b1) > (2^{15}-1) \text{ THEN } (2^{15}-1)$ $\text{ELSEIF } (a1-b1) < -2^{15} \text{ THEN } (-2^{15})$ $\text{ELSE } (a1-b1);$ $t2 = \text{IF } (a2-b2) > (2^{15}-1) \text{ THEN } (2^{15}-1)$ $\text{ELSEIF } (a2-b2) < -2^{15} \text{ THEN } (-2^{15})$ $\text{ELSE } (a2-b2);$
HSUB.us <i>ra,rb,rt</i>	$t1 = \text{IF } (a1-b1) > (2^{16}-1) \text{ THEN } (2^{16}-1)$ $\text{ELSEIF } (a1-b1) < 0 \text{ THEN } 0$ $\text{ELSE } (a1-b1);$ $t2 = \text{IF } (a2-b2) > (2^{16}-1) \text{ THEN } (2^{16}-1)$ $\text{ELSEIF } (a2-b2) < 0 \text{ THEN } 0$ $\text{ELSE } (a2-b2);$
HAVE <i>ra,rb,rt</i>	$t1 = (a1+b1)/2;$ $t2 = (a2+b2)/2;$
HSLkADD <i>ra,k,rb,rt</i>	$t1 = (a1 \ll k) + b1;$ $t2 = (a2 \ll k) + b2;$ (for $k = 1, 2, \text{ or } 3$)
HSRkADD <i>ra,k,rb,rt</i>	$t1 = (a1 \gg k) + b1;$ $t2 = (a2 \gg k) + b2;$ (for $k = 1, 2, \text{ or } 3$)

ss = signed saturation option
us = unsigned saturation

shifted in on the left and unbiased rounding* is performed on the least-significant bit on the right. Because the carry is shifted in, no overflow can occur in the HAVE instruction.

HSLkADD, or halfword shift left and add, allows one operand to be shifted left by *k* bits (where *k* is 1, 2, or 3) before being added to the other operand.

HSRkADD, or halfword shift right and add, allows one operand to be shifted right by *k* bits (where *k* is 1, 2, or 3), before being added to the other operand.

Both HSLkADD and HSRkADD use signed saturation.

* Unbiased rounding means that the net difference between the true averages and the averages obtained after unbiased rounding is zero if the results are equally distributed in the result range.

Saturation Arithmetic. In saturation arithmetic a result is said to have a positive overflow if it is larger than the largest value in the defined range of the result. It is said to have a negative overflow if it is smaller than the smallest value in the defined range of the result. If the saturation option is used for the HADD and HSUB instructions, the result is clipped to the maximum value in its defined range if positive overflow occurs and to the minimum value in its defined range if negative overflow occurs. This further speeds up the processing because it replaces using about ten instructions to check for positive and negative overflows and performs the desired clipping of the result for a pair of operations in one instruction.

Saturation arithmetic is highly desirable in dealing with pixel values, which often represent hues or color intensities. It is undesirable to perform the normal modulo arithmetic in which overflows wrap around from the largest value to the smallest value and vice versa. For example, in 8-bit pixels, if 0 represents black and 255 represents white, a result of 256 should not change a white pixel into a black one, as would occur with modulo arithmetic. In saturation arithmetic, a result of 256 would be clipped to 255.

Effect on MPEG Decoding. These parallel subword arithmetic operations significantly speed up several critical parts of the MPEG decoder program, especially in the inverse discrete cosine transform and motion compensation steps. More than half of the instructions executed for the inverse transform step are these parallel subword arithmetic instructions. Their implementation does not impact the processor's cycle time, and adds less than 0.2% of silicon area to the PA 7100LC processor chip. Actually, the area used was mostly empty space around the ALU, so that these multimedia enhancements can be said to have contributed to more efficient area utilization, rather than adding incremental chip area. See "Overview of the Implementation of the Multimedia Enhancements" on page 66.

Since the PA 7100LC processor has two integer ALUs, we essentially have a parallelism of four halfword operations per cycle. This gives a speedup of four times, in places where the superscalar ALUs can be used in parallel. Because of the built-in saturation arithmetic option, speedup of certain pieces of code is even greater.

System Optimization

The second longest functional step (see Table I) in MPEG decompression was the display step. Here, we leveraged the graphics subsystem to implement the color conversion step together with the color recovery already being done in the graphics chip.⁷ Color conversion converts between color representations in the YCbCr color space and the RGB color space. Color recovery reproduces 24-bit RGB color that has been color compressed into 8 bits before being displayed. Color compression allows the use of 8-bit frame buffers in low-cost workstations to achieve almost the color dynamics of 24-bit frame buffers. This leveraging of low-level pixel manipulations close to the frame buffer between the graphics and video streams also contributed significantly to the attainment of real-time MPEG decompression. Color recovery and the graphics chip are described in the articles on pages 51 and 43, respectively.

Other PA 7100LC processor enhancements streamline the memory-to-I/O path. By having the memory controller and the I/O interface controller integrated in the PA 7100LC chip, overhead in the memory-to-frame-buffer bandwidth is reduced. Overhead in the processor-to-graphics-controller-chip path is also reduced for both control and data.

Path Length Reduction

Table III shows the same information as Table I but for the low-end Model 712 workstation which uses the multimedia-enhanced PA 7100LC processor and the graphics chip mentioned above.

Table III
Instructions and Time Spent in each MPEG Decompression Step on a Model 712 Workstation

	Millions of Instructions	Path Length (%)	Time (%)
Header decode	0.60	0.2	0.3
Huffman decode	55.0	16.1	14.5
Inverse quantization	8.9	2.6	4.5
Inverse discrete cosine transform	138.5	40.6	34.4
Motion compensation	74.8	21.9	25.6
Display	63.0	18.5	20.7
Total	340.8	100.0	100.0

The Model 712 executes consistently fewer instructions than the Model 735 for the same MPEG decompression of the same video clip. It is also faster in MPEG decompression even though it operates at only 60% of the 99-MHz rate of the high-end Model 735 and has only one eighth of the cache size. This shows the performance benefits from the path length reduction enabled by the PA-RISC processor and system enhancements for multimedia acceleration.

Performance

The performance of the PA-RISC architectural enhancements and the leveraging of the graphics subsystem for video decompression can be seen in Fig. 5. This data is for a




Fig. 5. Maximum MPEG decode frame rates for different models of HP 9000 Series 700 workstations. These rates are for a 352-by-240-pixel clip that was encoded at 30 frames/s.