VOLUME 10, NUMBER 16
DECEMBER 9, 1996

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Mitsubishi Designs DVD Decoder

D30V's LIW Core Processes Audio, Video in Consumer Products

by Linley Gwennap

Joining the crowd of vendors chasing the emerging DVD market, Mitsubishi has developed the D30V, a programmable media processor intended mainly for consumer devices. By eschewing the floating-point and interface logic needed for PCs, the company has achieved a smaller die size than any announced PC media processor: just 37 mm^2. This tiny die should give the D30V an advantage in the cost-conscious consumer-products market.

Taking a page from its Alpha processors, Mitsubishi set the target clock speed for the D30V ahead of any other media processor's: 250 MHz. The high clock speed allows the chip to use a relatively simple two-way LIW (long instruction word) architecture rather than the more complex VLIW (very long instruction word) design in the Philips Trimedia processor. The D30V's large register file, guarded execution, and optional serialized execution foreshadow techniques likely to be included in next-generation general-purpose instruction sets such as the Intel/HP IA-64.

To speed processing, the chip contains 64K of on-chip memory. It takes in a full DVD bit stream and produces standard video and audio output. The D30V core can perform system-level functions as well as media processing, eliminating the need for an external host processor in most applications. The chip recently attained first silicon and is expected to appear in DVD players and other devices late next year.

LIW Core Contains Two Function Units

Each D30V instruction is 64 bits wide and generally contains two subinstructions. These subinstructions correspond to the two main function units in the processor core, shown in Figure 1. Both function units have 32-bit data paths. The memory unit handles all accesses to main memory (loads and stores). If a memory access is not required in a particular cycle, this unit can perform basic integer arithmetic and shifts. The memory unit also handles all branch instructions.

The second unit is the main integer unit. Like the memory unit, it handles basic integer arithmetic and shifts. The integer unit, however, adds a 32 × 32-bit multiplier as well as two special accumulators for multiply-accumulate (MAC) operations.

Both function units share access to a multiported register file. This file contains 64 registers of 32 bits each, twice as many registers as in a standard RISC processor. The additional registers come in handy when unrolling loops and for storing lots of matrix coefficients. Several instructions can access the register file as 128 entries of 16 bits each, further increasing its effective capacity.

SIMD Operations Improve Parallelism

The D30V contains a set of instructions that operate on both halves of a 32-bit register independently, as in Intel's MMX and other multimedia extensions. These SIMD (single-instruction multiple-data) instructions include add, subtract, multiply, and shift instructions.

The D30V can perform two 16 × 16-bit multiplies in parallel using both halves of the source registers; the products can be truncated to 16 bits and stored in the same register, or the full 32-bit products can be stored in two

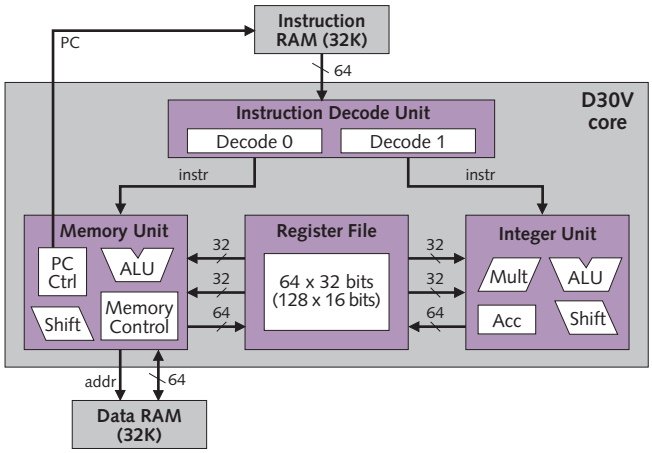

Figure 1. The D30V core has an LIW instruction decoder, two main function units, and a 64-entry register file.

Mitsubishi D30V

consecutive registers. The 32 × 32-bit multiply also has two options. If the product is known to be only 32 bits, it can be stored in a standard register. For larger values, the product must be stored in one of the two 64-bit accumulators. These accumulators are also used by the MAC operation, which adds the operation's 64-bit product to the contents of the accumulator. The value in the accumulator can later be transferred to the general registers.

The 32-bit multiply and MAC operations take two cycles to complete but are fully pipelined, so a new operation can be launched every cycle. Multiplication of 16 bits can be completed in a single cycle.

Unusual operations include AVG, which calculates the average of two values, a task that otherwise requires an add and a shift (divide by two). This instruction is also found in HP's PA-RISC multimedia extensions.

The chip also performs an ADDS (add sign) operation that adds the sign bit of one register to the contents of another. This operation is used to implement integer division with truncation toward zero, a function required in the inverse quantization phase of MPEG decoding.

Mitsubishi's SAT operation allows signed and unsigned saturation to an arbitrary bit position; most processors can saturate only to 8-, 16-, or 32-bit boundaries. Because many multimedia data types have widths such as 12 or 20 bits, the SAT operation can be more efficient for these calculations. Fujitsu's MMA media processor (see **101502.PDF**) has a similar capability.

Unlike many media processors, the D30V does not currently implement saturating arithmetic (other than the SAT operation) and doesn't perform SIMD operations on 8-bit data. Mitsubishi's technique is to promote 8-bit data to 16-bit halfwords. This promotion allows up to 8 bits to capture any overflows or underflows. In some algorithms, the extra bits provide greater precision. After the calculations are complete, only a single SAT operation is needed.

Although other processors are able to achieve a higher degree of parallelism in some restricted cases by working on 8-bit data with saturating arithmetic, for the video algorithms that the D30V is focused on, the extra precision of 16-bit operations is needed. Thus, the company saw little value in adding support for 8-bit arithmetic. Note that Chromatic's Mpact achieves greater parallelism than the D30V, even on 16-bit operands, using its wider (72-bit) data path.

	Fetch	Decode	Execute	Writeback
ALU or Mul16	Fetch instr	Decode instr	ALU op	Store result
Load or Store	Fetch instr	Decode instr and calc addr	Mem access	Store result
Branch	Fetch instr	Decode instr and calc target	Branch	

Figure 2. By calculating the memory address in the decode stage, the simple four-stage pipeline in the D30V eliminates the load-use penalty. One delay slot is needed to alleviate the branch penalty.

Flexible Addressing for 16-Bit Data

The D30V goes beyond most media processors in optionally accessing the register file in 16-bit quantities. Most of the instructions mentioned above can operate on any 16-bit halfword in the register file. Although this method halves throughput compared with the SIMD operations, it provides more flexibility in handling special cases.

In this mode, the high portion of the register is addressed using an "H" suffix, while the low portion uses an "L" suffix. Using this notation, an ADD operation can perform the function r2H + r3L → r1L, for example, without disturbing the contents of the upper half of r1. By using r0 (which always reads as zero) as one of the operands, this operation can copy data from one half-register to another.

To pack and unpack halfwords, Mitsubishi provides a JOIN operation that concatenates any two 16-bit values from the register file into a single 32-bit register. This operation can also be used to rearrange data, a function required in the IDCT (inverse discrete cosine transform) portion of MPEG-2 video decoding and in other multimedia algorithms.

Using a single subinstruction, the chip can load four consecutive bytes and unpack them into four consecutive half-registers using signed and unsigned formats. This operation converts 8-bit data to the 16-bit format used by the D30V's arithmetic instructions.

In the converse case, a single operation can combine data from four consecutive half-registers, storing a single 32-bit value to memory. In this case, the upper byte of each half-register is discarded. Other variations store halfwords from two registers or load two halfwords into two 32-bit registers. The memory unit has a 64-bit path to the register file to facilitate these operations. Together, these operations allow efficient loads and stores for up to four operands per cycle. The D30V also includes simple 32-bit and 64-bit loads and stores.

Any memory access can choose from several addressing modes. The chip includes the standard register+immediate and register+register addressing modes, adding an option to increment or decrement the base register after the access. This option is handy for loops, eliminating the need for a separate increment operation. Like many DSPs, the D30V supports a modulo addressing mode. This mode is good for accessing circular buffers and queues of various sizes.

Branches Have Variable Delay Slots

Despite the high clock speed, the D30V requires only four pipeline stages, as Figure 2 shows. The LIW architecture eliminates the delays of on-chip pairing logic. In fact, instruction decoding is so simple that Mitsubishi is able to pull the address calculation for memory or branch instructions into the second (decode) stage. This change allows memory

accesses to occur in the third (execute) stage rather than the fourth stage, as in the traditional RISC pipeline. This in turn eliminates the load-use penalty.

In fact, the pipeline has very few hazards. All instructions execute in a single cycle except for MUL32, which takes two cycles. There is a single interlock cycle if the next instruction requires the result of the multiply, which can be easily avoided with good coding practices. Bypass logic avoids most other potential interlocks.

Branches are the main cause of pipeline disruption. For simplicity, the D30V has no branch prediction. Because the target address of a branch is not known until the end of stage two, a standard branch instruction has a one-cycle penalty in the event of a taken branch.

Like many processors, the D30V solves this problem with a delayed branch. With a simple delayed branch, the subsequent instruction is always executed, so the penalty cycle is avoided. Mitsubishi allows branches to specify up to 64 delay slots. A programmer can place the branch operation in any convenient location that might otherwise have been an unused instruction slot, increasing the efficiency of the operation pairing. This method is particularly useful for unconditional branches such as subroutine calls and returns, since there is no need to wait for the condition to be calculated.

Conditional branches are very restricted: the only conditions are whether a register is zero or nonzero. More complicated conditions must be computed using explicit compare instructions.

Like many DSP chips, the D30V has a zero-overhead REPEAT instruction. This operation creates a tight loop of several instructions that are executed up to 2^{32} times. The repeat operation specifies the number of instructions in the loop and the number of iterations.

Pairing Encoded in Instructions

As Figure 3 shows, the basic 64-bit instruction format specifies two subinstructions, the left and the right. Most of the time, the left subinstruction is issued to the memory unit while the right subinstruction goes to the integer unit, both at the same time. Two bits, FM[0:1], can specify three other modes of operation that are occasionally useful.

In some situations, two consecutive subinstructions cannot be executed in the same cycle due to data dependencies or resource conflicts. Because the hardware doesn't check for such dependencies, the compiler (or programmer) must recognize this situation. Of course, the best solution is to rearrange the code to avoid the problem, but if this is not possible, the instruction can be encoded to execute the two subinstructions serially. This encoding avoids the need to bloat the code by inserting NOPs.

If FM[0:1]=01, the left subinstruction is executed first, followed by the right subinstruction on the following cycle. In this case, the left subinstruction must be issued to the memory unit, simplifying the issue logic; the right subinstruction, which has an extra cycle to be decoded, can be either a memory or integer operation. In the case where an integer subinstruction must be executed first, it is placed on the right side and FM[0:1] is set to 10.

Finally, the D30V supports a few operations that take up the entire 64-bit instruction word. These operations include arithmetic, memory, and branch operations with full 32-bit immediate values. Since most processors require two 32-bit instructions to form a 32-bit immediate value, this encoding requires the same space as in other processors and avoids sacrificing a register for the long immediate value.

For More Information

Mitsubishi has not yet announced the price or availability of the D30V. We expect the chip to be available in 2H97. For more information, contact Mitsubishi Electronics America (Sunnyvale, Calif.) at 408.730.5900. There is no information yet available on the Web.

Guarded Execution Eliminates Branches

For each subinstruction, the instruction encoding has three guard bits that specify any Boolean combination of the two user flags in the PSW (program status word). The subinstruction is executed only if the specified condition is met.

This technique, which is also used in the Trimedia processor (see **091506.PDF**), eliminates branches in many cases. Instead of branching to two different routines to handle a particular value, for example, the code from the two routines could be in line, with each routine guarded by the appropriate condition. Assuming the routines are short, this method would give the compiler more flexibility in scheduling operations, since operations from one routine could be mixed with operations from the other routine or with operations from the in-line code.

This technique can also simplify the handling of delayed branches. Some or all of the operations in the delay

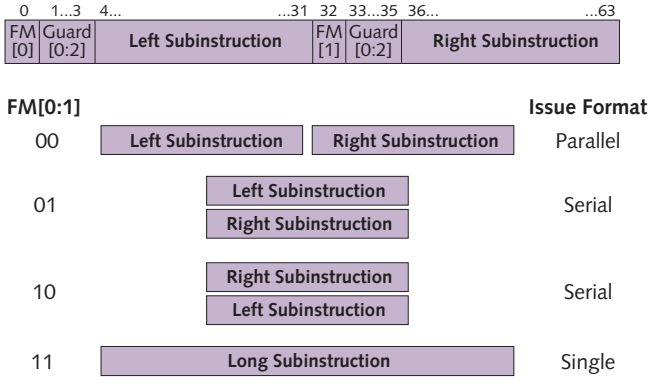

Figure 3. The 64-bit instruction word supplies two standard subinstructions or one long subinstruction. Depending on the state of the FM bits, the subinstructions can be executed together or one at a time.

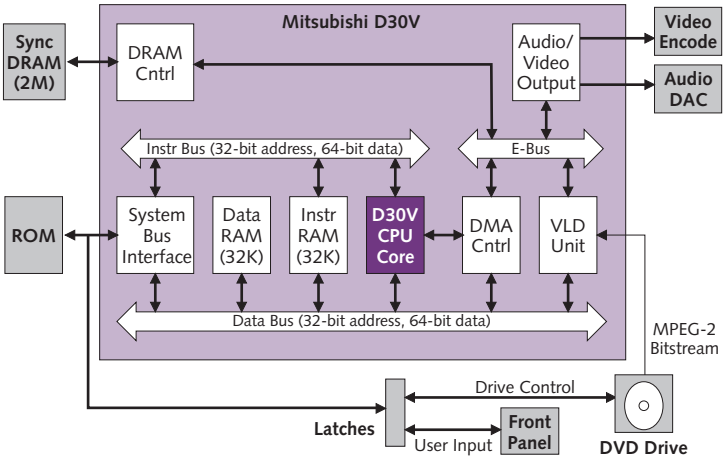

Figure 4. Few external components must be added to the D30V to form a complete DVD player.

slots can be guarded with the branch condition, so they are executed only if the branch is taken. This technique makes it easier to fill the delay slots. Guarding is critical in the Trimedia architecture, which requires three branch delay slots that each specify five operations. For the D30V, which has a minimum of one delay slot with two operations, this feature provides a smaller but important performance boost.

After accounting for the FM and guard bits, each subinstruction has only 28 bits instead of the 32 used for typical RISC instructions. The larger register file requires six bits instead of five to specify each of up to three operand registers. These factors reduce the number of bits left for opcode extensions, branch displacements, and immediate values. For example, most branches have only a 12-bit displacement; if a larger displacement is required, the long instruction format must be used.

Few External Components Needed

The flexibility of the D30V instruction set allows most of the DVD decoding process to occur using only the programmable core. The only special-purpose hardware on the chip, as Figure 4 shows, is a variable-length decoder (VLD) that handles the first level of decompression for the MPEG bit stream. The chip also contains 32K of instruction memory, 32K of data memory, a DMA engine, and interfaces to audio and video output devices, external DRAM, and ROM. The on-chip memory is implemented as SRAM, not cache, so software must decide which items are kept in the fast local memory instead of the slower external DRAM.

For a typical DVD player, few external components are needed. The DVD drive must supply a bit stream to the D30V. The Mitsubishi chip requires 2M of synchronous DRAM, a ROM to hold the LIW code, and audio and video DACs. No host processor is required; the D30V can perform functions such as interpreting input from the control panel and selecting the appropriate track. The initial version does not provide parallel I/O for this purpose, but inexpensive latches can be added to the system bus, as Figure 4 shows, to connect to the user interface. Future versions may integrate more of this I/O connectivity.

This simple design reduces system cost compared with other media processors. The integration of the D30V is best matched by the Fujitsu MMA, but that chip requires an external SparcLite host processor. Neither Samsung's MSP-1 (see **101101.PDF**), which includes an ARM core in addition to its vector processor, nor Philips' Trimedia chip requires an external host CPU.

At 37 mm^2, the D30V is less than half the size of the MSP-1 or Trimedia chips, due in part to Samsung's and Philips' decision to target PC applications by including floating-point logic and PCI support. Mitsubishi should have a clear manufacturing-cost advantage over these chips for DVD players and similar devices.

DVD Decoding with Cycles to Spare

Based on Mitsubishi's simulations, a 250-MHz D30V should handle full MPEG-2 decoding (audio and video) using about 90% of the total compute bandwidth of the processor. That leaves about 10% of the processor's power for the user interface and other host-processor tasks.

On inner loops such as IDCT and motion estimation, both issue slots are occupied on 100% of the instructions, according to Mitsubishi. The flexible nature of the two function units, along with the variable-delay branches, helps make this efficiency possible in hand-coded loops. This efficiency keeps the number of cycles for a complete IDCT to 450, or 1.8 microseconds.

The company also noted that 90% of the operations in the motion-compensation loop take advantage of halfword operations, all using the SIMD instructions. IDCT, on the other hand, performs just 17% of its operations on halfwords, and about two-thirds of these operations calculate only one halfword per cycle rather than two.

Although the D30V's primary target is DVD players, it can be used for other multimedia applications as well. The 250-MHz chip can complete a 256-point complex radix-2 FFT, for example, in just 40 µs. The FFT is frequently used in

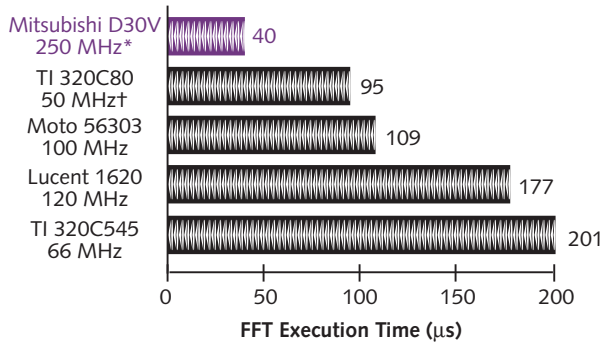

Figure 5. The D30V outperforms common DSP chips on the 256-point complex FFT benchmark. †using one of four on-chip processors (Source: Berkeley Design Technology except *Mitsubishi)

audio and communications algorithms. By comparison, a 100-MHz 56303 from Motorola, one of the fastest DSP chips available, requires 109 µs for the same task, as Figure 5 shows. Thus, the D30V has enough horsepower to do tasks handled today by conventional DSP chips, although these chips are likely to be less expensive for such relatively simple tasks.

The D30V will be fabricated in 0.3-micron four-layer-metal CMOS. In this advanced manufacturing process, the 300,000-transistor LIW core requires just 8 mm^2, and the entire processor, including 64K of on-chip SRAM and the on-chip system logic, consumes 37 mm^2. This makes the D30V by far the smallest media processor yet disclosed; the closest competitor is Fujitsu's MMA at 77 mm^2.

The Mitsubishi chip will use an inexpensive plastic package, implying a low power dissipation. The company would not, however, provide details on the pin count or power dissipation of the part.

More Powerful Than Competitors

Many media processors are straddling the fence between the PC and consumer markets, hoping to gain design wins in whichever camp grows faster. Mitsubishi, one of the world's largest makers of consumer electronics, has chosen to focus only on consumer products. This focus should give the D30V a cost advantage over many competitors.

Among media processors aimed at the consumer market, Fujitsu's MMA appears to offer the strongest competition. MMA, like the D30V, maintains a small die size and low power consumption by eliminating support for 3D graphics and the PCI bus. Until Mitsubishi reveals the full details of its design, a precise comparison is impossible, but the two chips appear similar by these metrics.

The D30V, however, appears to offer a performance advantage over the Fujitsu chip. DVD decoding consumes the entire bandwidth of MMA, leaving no cycles for other functions. Thus, MMA relies on an external host processor for the user interface. Based on its specifications, the D30V is 40% faster than MMA in clock speed, has more flexible instruction-pairing rules, and has four times as much on-chip SRAM. These advantages help it perform DVD decoding with cycles to spare, eliminating the need for an external host processor.

The Fujitsu chip appears to have a time-to-market advantage, however, as it is due to reach production in 1H97, about six months before the D30V. Fujitsu also offers a full set of development tools from Green Hills for its MMA; Mitsubishi is currently developing a set of Gnu tools, with availability set for 2Q97.

Tough competition will also come from hardwired devices. Mitsubishi itself offers hardwired MPEG decoders, and all first-generation DVD players are based on such solutions. Oak (see **1015MSB.PDF**) and others are considering combining hardwired audio and video decoders on a single chip. Mitsubishi will have to price the D30V aggressively to beat these designs.

D10V Aims at Cell Phones

Mitsubishi's D10V uses a slimmer version of the LIW core found in the D30V. The smaller chip's instructions are only 32 bits wide but still contain two subinstructions; each is 15 bits wide. There is no guarded execution, and the small instructions can specify only two operands instead of three. As with the D30V, a "long" instruction format combines the two halves, in this case into a single 30-bit subinstruction.

The register file is 16×16 bits instead of 64×32. The smaller registers can't support the SIMD instructions found in the D30V. The multiplier is physically smaller, as no 32×32 capability is needed, and the accumulator registers are only 40 bits wide. Otherwise, the D10V has the same number and type of function units and same basic design as its big brother.

The D10V is aimed at audio and communications applications, while the D30V handles video applications. With the lower performance demands for audio, the D10V clocks at a sedate 52 MHz. With one-fifth the clock speed and a simpler design than the D30V, the D10V consumes less than 100 mW. Mitsubishi claims that, even at this clock rate, the small chip can perform all required signal processing for a GSM cellular phone with just 30% of its total CPU bandwidth.

The D10V is built in a 0.35-micron three-layer-metal CMOS process, slightly less aggressive than the D30V's process. Mitsubishi is now sampling the D10V; we expect it to reach production around mid-1997.

Good Match for Consumer Products

The drawback to tightly focusing on the DVD decoder market is that the D30V is suited for few other applications. The video capabilities of the chip may come in handy in a videophone or other video-conferencing device. But without floating-point capability and other special functions, the D30V won't fit well into a video-game player, although Mitsubishi touts this application. The chip is unsuited to PC applications for the same reason. The company may add these capabilities in a future version aimed at PCs.

While the D30V is of relatively low cost compared with other media processors, it is likely to be more expensive than standard DSP chips. Mitsubishi has developed the D10V (see sidebar) to address applications that process only audio data, such as modems and cell phones.

By carefully matching the capabilities of its chip to the requirements of the consumer market, Mitsubishi has created a powerful yet inexpensive design. Although it has missed the initial wave of DVD players, the chip should be ready in time for the 1997 holiday season. With Mitsubishi's own consumer division as a sure customer and many other potential buyers, the D30V looks to do well in this emerging market. M