

General Purpose Processors for Floating-Point DSP

"DSP" is an application space and the set of microprocessors which address the needs of those applications. When selecting a microprocessor there are four key attributes that facilitate DSP performance. First, the microprocessor must perform simple arithmetic operations every clock cycle. Second, it must be able get the data in and out of the processor in a quick and controllable manner. Third, it must be supported by a powerful development environment tailored to real-time applications. And fourth, for embedded DSP, it must optimize computational speed with considerations to size, power, and cost objectives.

The term "DSP" describes an application space and the set of microprocessors that have been designed to address the needs of those applications. When selecting a microprocessor for floating-point digital signal processing, there are four key attributes of a microprocessor that facilitate DSP performance. First, the microprocessor must perform a multiply-accumulate (MAC) and other simple arithmetic operations every clock cycle. Second, it must be able get the data in and out of the processor in a quick and controllable manner. Third, it must be supported by a powerful software development environment tailored to real-time applications. And fourth, for embedded DSP, it must not only optimize computational speed, but give equal weight to size, power, and cost objectives.

Traditionally, these four attributes were found only in processors specifically designed for DSP from companies like Texas Instruments, Motorola, and Analog Devices. However, general purpose microprocessors designed with the laptop and embedded markets in mind can also meet those requirements.

PERFORMANCE MEASUREMENT

The raw measurement of a processor's speed is the number of operations performed per second. This is simply the processor clock cycle multiplied by the number of operations that can be performed simultaneously each cycle. Using such peak performance numbers for evaluation can be misleading without an understanding of the underlying architecture. This is especially true when comparing two chips of different architecture classes, such as RISC and DSP. DSP chips historically implement a number of "arith-

metic" overhead functions in special purpose hardware that are not typically counted in the MOPS rating. Loop counters and address generators are two such overhead operations implemented in hardware with "zero overhead". RISC processors generally lack these special purpose hardware functions, and therefore such overhead operations must be performed by the general purpose ALUs. This makes the number of integer ALUs of great importance, even for floating-point applications.

GETTING MADD

The first requirement on a microprocessor for DSP applications is that it must perform a multiply-accumulate (MAC), or the more general multiply-add (MADD), and other simple arithmetic operations every clock cycle. A RISC chip can perform single cycle execution for simple arithmetic operations such as a MADD. Some DSP-specific architectures, like the TI C40, only execute a MAC every cycle by running the internal clock rate at half the speed of the external clock. The SHARC chip has an advantage here for FFT applications in that it can do a subtract along with the multi-



Figure 1. SHARC chip

ply and add all in the same 25 ns cycle. On the RISC side, a PowerPC 604e running at 333 MHz can execute a single-precision MAC on each 3 ns processor cycle. If double precision calculations are required, some RISC processors have a distinct advantage. The PowerPC 604e can execute a double-precision operation every cycle, which is the same speed as single-precision calculations. Although this extra precision is not typically necessary for DSP applications, the reduced precision that results from a very long string of operations, such as a long FFT, can require a 64-bit calculation just to get the accuracy of a 32-bit floating-point number.

Higher floating-point throughput can be achieved using some of the newer architectures. (Very Long Instruction Word) VLIW and SIMD (Single Instruction Multiple Data) architectures use multiple execution units to execute multiple parts of a complex instruction or multiple instances of the same instruction to achieve higher processor throughput. These architectures can be found in both DSP and RISC processors. For example, both the SHARC 21160 and PowerPC G4 with Altivec are SIMD architectures. A notable difference is that multiple execution units in a DSP chip are typically not capable of each performing a MADD. One unit does the multiply and another does the add. Conversely, each floating-point execution unit in a RISC chip typically can perform a MADD.

MEMORY ARCHITECTURE ISSUES

Most DSP applications are memory bandwidth limited. The traditional DSP architecture, the Harvard architecture, addresses this crucial architectural attribute by employing dual memory buses. For example, the TI C40 excels at I/O because it provides two 32-bit data paths to memory, allowing for two 32-bit numbers to be read from memory on every cycle. However, writes take two cycles and become the limiting factor for throughput in many applications. The SHARC (Super Harvard Architecture) employs two buses internally but only one 48-bit data path to external memory, which allows one 32-bit floating-point number to be written each cycle. Reads by the SHARC typically take 2 cycles from external memory. To compensate, the SHARC includes SRAM in the chip and allows for more efficient data movement off the chip by utilizing its internal DMA engine to copy data from its two internal SRAM banks to external memory. However, making effective use of these architectures, that is utilizing twin external data buses or DMA transfers from twin internal memory banks, can be a tricky operation. The details of this are best applied to applications with a small




Figure 2. Excalibur card

amount of code, typically centered around a particular inner loop.

By comparison, RISC chips including the PowerPC, have a single 64-bit data path to external memory. The PowerPC 604e has an external memory access rate of 83.3 MHz. This 667 MB/sec bandwidth is more than four times the bandwidth of the available DSP solutions. Having fewer but wider memory paths also allows for a significantly faster memory design and a simpler programming model. Thus, RISC processors are more applicable to a wider range of applications, beyond those containing just a tight inner loop.

Getting the most out of the external memory interface of any processor chip can be a challenge. Even when using fast SDRAM in burst mode, the access latency is great enough to lower the overall sustained throughput on reads to a fraction of the peak throughput.

Many DSP chips do not support pipelined memory accesses. Some architectures support pipelining with a separate DMA engine. The PowerPC 740/750 supports only limited pipelining, optimizing it for alternating instruction and data accesses. This is a simplified design that can work well on general purpose applications that have an unstructured mix of instruction and data transfers, but it is not of much benefit for signal processing applications. More flexible data pipelining, such as that found on the PowerPC 604e, provides the most benefit to signal processing applications.

SOFTWARE ARCHITECTURE ISSUES

One of the most important factors for selecting a DSP chip is the quality and ease-of-use of the available software tools and development environment. Due in large part to the limited market size, pure DSP chips typically have fewer software tools available than do RISC

PROCESSOR

chips. A RISC architecture that runs laptop and desktop computers, parallel supercomputers, and VME boards has broader software support.

As military usage of DSP boards increases, the need to develop large amounts of code that is maintainable, and to provide for processor upgrades that run the same code is driving program managers to use high-level languages even for DSP portions of the application. RISC chips like the PowerPC support a high level of performance with software portability because these processors are designed to support HLL compilers. Some compilers for the PowerPC, like the SKYvec compilers from SKY Computers, automatically strip mine vectors, chain vector functions, and eliminate temporary vectors.

THE RIGHT TOOL FOR THE RIGHT JOB

Although the PowerPC's architectural advantages are especially useful for large systems, the typical embedded parameters of MFLOPS/Watt, MFLOPS/\$, and MFLOPS/sq. in. also need to be factored in. To account for the various memory and support chips needed, this comparison is best done at the board level as shown in table 1.

TI chips, for instance, seem to have found a home in computer telephony because of their low cost and good all-around DSP performance. The SHARC excels at FFTs that fit into the on-chip SRAM, but performance drops off greatly if off-chip accesses are required. This combined with the high density at which you can populate a board with these chips, makes the SHARC an excellent choice for front-end signal processing of limited sized data sets. For larger data set sizes and good performance beyond the tight inner loop, the PowerPC 604e is the best choice.

Because the choice of processor is dependent on the algorithms being run, SKY Computers offers both DSP and RISC-based solutions for DSP applications. The SHARC chip is often a good match for front-end signal processing, for which SKY sells a 12-processor SHARC-based card called SHARCpool. Particular attention was paid to the design of the memory architecture of this card, and the result is the ability to transfer data in the shared memory of a SHARC cluster at 320 MB/sec. For back-end processing, SKY offers 4 PowerPC 604e processors on the Excalibur card. With 2.66 GFLOPS and up to 512 MB of SDRAM accessed at 667 MB/sec by each processor, Excalibur is also a good choice for front-end processing when

	INTEL I860	TI C40	ADI 21062	POWERPC 604E
Processor Specs				
CPU Clock	40 MHz	50 MHz	40 MHz	333 MHz
Memory Bus	160 MB/sec	2 * 200 MB/sec	160 MB/sec	667 MB/sec
On-chip Memory	4 KB Instr cache 8 KB Data cache	0.5 KB Instr cache 8 KB RAM	256 KB RAM	32 KB Instr cache 32 KB Data cache
Processor Benchmarks				
1K CFFT	0.75 ms	1.28 ms	0.46 ms	0.23 ms
1K RFIR (64 taps)	1.76 ms	1.48 ms	1.79 ms	0.32 ms
Typical 6U VME Implementations				
Processors per board	4	8	12	4
MFLOPS per board (single precision)	320 (240 double precision)	400	1440 (for FFTs) 960 (non-FFT)	2666
\$/MFLOPS on a board	\$78	\$54	\$18 (FFT) \$27 (non-FFT)	\$16
Software Tradeoffs				
Software Tools	Mature, High-level, Intelligent	Mature, Low-level	Immature but growing rapidly	Mature, High-level, Intelligent
Effective Algorithms	Complex DSP algorithms with off-chip accesses	Small embedded DSP applications (e.g. computer telephony)	On-chip FFTs for front-end sonar, radar and signal processing	Complex DSP and data processing algorithms with large data sets

Table 1. Board level comparison

INITIAL OFFERING FOR	ADI 21160	TI C67	MOTOROLA G4 WITH ALTIVEC
CPU			
CPU clock	100 MHz	167 MHz	≥ 350 MHz
FP Instr per cycle	6	6	4 + 1
Max OPS per Instr	1	1	2 (MADD or MSUB)
MFLOPS (peak)	600	1000	> 2800
Memory			
On Chip Memory/Cache	256 kbytes	128 kbytes Instr 128 kbytes Data	≥ 32 kbytes Instr ≥ 32 kbytes Data
External Interface	64 bits @ 50 MHz	32 bits @ 167 MHz	64-bits @ 100 MHz
Peak Memory Bandwidth	400 MB/sec	667 MB/sec	800 MB/sec
Memory Pipeline Depth	1	1	≥ 2
Real World			
1K CFFT with bit reversal	90 μs	108 μs	30 μs
Power (typical)	2 W	2 W	< 10 W
Processors per 6U VME	8	4	4
GFLOPS per 6U VME	4.8	4.0	≥ 11

Table 2: Future floating-point DSP and RISC offerings

the task is complex so that the ease-of-use and extra flexibility of a RISC chip become more important. Either of these cards can be attached to a 6U VME board, and four of them can be attached to a 9U VME board. Applications can mix-and-match SHARCs and PowerPCs in the same system or even on the same motherboard in the case of the 9U solution.

FUTURE DIVERGENCE OR CONVERGENCE?

As DSP chip vendors search for ways to improve performance beyond standard Harvard architectures, they have turned in several different directions, including superscalar, VLIW, and SIMD architectures. For example, the TI 'C67 provides extra floating-point and fixed-point units access through a VLIW architecture. The new ADI 21160 SHARC processor uses two groups of floating-point execution units in a SIMD model to double the number of operation per cycle in addition to doubling the clock frequency.

While these differing strategies appear to signal a divergence in the DSP chip industry, this can also be looked at as a convergence with the direction of RISC processor architecture. This convergence includes specific instances, such as the design of the TI 'C67 to depend on the extra execution units instead of special hardware to perform the DSP overhead functions – just like a RISC chip. Such design choices also have broad implications, such as the increased difficulty to program in assembly language, so the programming model has also moved closer to that of a RISC chip.

From the other side, RISC processors targeted at multimedia applications have included DSP cores. In a

more integrated approach, the PowerPC G4 incorporates a vector processor along side of the standard ALUs and floating-point unit. This vector unit follows a SIMD programming model, as does the ADI 21160. All of these processors are operating clocks well above the typical DSP clock frequency of 40–60 MHz. Based on the FFT execution times of these parts projected by the respective chip manufacturers (shown in table 2), the PowerPC G4 appears to have the jump on the super-charged DSP chips.

Choosing the right processor is ultimately dependent upon the requirements of the application. As a general rule, the SHARC processor fits best in front-end signal processing applications. A RISC processor with a fast memory interface, such as the PowerPC 640e, is the chip of choice for complex algorithms or mixed signal and data processing. ■

Stephen Paavola is a product manager at SKY Computers, Inc. of Chelmsford, MA. He has over 25 years experience in senior engineering, sales and marketing roles at Digital Equipment Corporation, Cetia, and SKY. Mr. Paavola holds a BS from Cal Tech.