

Benchmark Evaluations of Modern Multi-Processor VLSI DSP μ Ps

Aaron L. Robinson and Fred O. Simons, Jr.
High-Performance Computing and Simulation (HCS) Research Laboratory
Electrical Engineering Department
Florida A&M University and Florida State University
Tallahassee, FL 32316-2175

Abstract - The authors continue their tradition of presenting technical reviews and performance comparisons of the newest multi-processor VLSI DSP μ Ps with the intention of providing concise focused analyses designed to help established or aspiring DSP analysts evaluate the applicability of new DSP technology to their specific applications. As in the past, the Analog Devices SHARCTM and Texas Instruments TMS320C80 families of DSP μ Ps will be the focus of our presentation because these manufacturers continue to push the envelop of new DSP μ Ps (Digital Signal Processing microProcessors) development. However, in addition to the standard performance analyses and benchmark evaluations, the authors will present a new image-processing bench marking technique designed specifically for evaluating new DSP μ P image processing capabilities.

1. Introduction

The combination of constantly evolving DSP algorithm development and continually advancing DSP μ P hardware have formed the basis for an exponential growth in DSP applications. The increased market demand for these applications and their required hardware has resulted in the production of DSP μ Ps with very powerful components capable of performing a wide range of complicated operations. Due to the complicated architectures of these new processors, it would be time consuming for even the experienced DSP analysts to review and evaluate these new DSP μ Ps. Furthermore, the inexperienced DSP analysts would find it even more time consuming, and possibly very difficult to appreciate the significance and opportunities for these new components. Thus, performance comparisons of modern multiprocessor VLSI DSP μ Ps will be presented in hopes of providing brief time saving reviews for analysts who need to consider these devices for critical new applications. The performance comparisons will emphasize the Analog Devices SHARCTM and the Texas Instruments TMS320C80 and TMS320C60 families. The Motorola DSP96002 will also be included because of its data precision.

Initially, distinguishing architectural features of each of the microprocessors are discussed. Common processor features are then compared and presented in tabular form. The processing speed and power of each of the DSP μ Ps will then be evaluated in terms of some common digital filtering operations and simple image processing applications. The DSP μ Ps performance characteristics are presented in *easy-to-interpret* graphical forms.

It should be noted that the digital filtering operations presented were not chosen as a basis for performance criteria because of their implementation and computational complexities. Instead, the operations were chosen due to the fact that they characterize performance criteria common to almost all DSP applications; e.g., implementation algorithm and data multiplication, addition, quantization, etc. A similar line of logic was applied to the choice of image processing applications. The applications chosen are intended to illustrate a few simple and basic operations involved in point, spatial, and frequency domain image processing while also emphasizing the features of processors specifically adapted to these operations.

2. Modern Parallel DSP μ P Candidates and Features

The Texas Instruments TMS320C6201 is a fixed point DSP μ P based on the VelociTI 'C6200 CPU. The processor is capable of 5ns-instruction cycle @ 200 MHz. The architecture features two 16-Bit Multipliers (32-Bit Results), six Arithmetic Logic Units and 1 M-bit of on-chip RAM. In addition to the eight independent functional units, there are 32 32-Bit general purpose registers. The 256-bit wide program memory is capable of fetching eight 32-Bit instructions every cycle. Processor operation at eight instructions per cycle is approximately 10 times faster than other available processors.

The TMS320C6201 also features a 32-Bit external interface unit (EMIF). The EMIF allows for glueless interfacing to synchronous and asynchronous memories and other external controller devices. The device also has a 2 channel Direct Memory Access Coprocessor and a 16-Bit host access port for processor to processor communication.

The Analog Devices ADSP-21060 SHARC™ (Super Harvard Architecture Computer) is a 32-bit processor that builds on the ADSP-21000 core to form a complete system on a chip. It integrates the ADSP-21020, the fastest 32-bit IEEE floating point DSP, with 4 megabits of on chip SRAM to form a powerful computational system. The ADSP-21060 on chip memory represents the largest on chip memory available on any current DSPuP. The processor's on chip DMA controller along with its dual ported memory enables efficient data and

instructions transfers by allowing simultaneous memory accesses by both the DMA controller and the core processor.

The ADSP-21060 serial and parallel communications ports combine with the host port and multiprocessor interface to increase the connectivity of the SHARC™ and therefore make it better suited for multiprocessing. In addition to the host port and multiprocessor interface, the ADSP-21060 also features six 4-bit link ports that provide additional I/O and multiprocessing capabilities. Because the bit capacity of the ports is small, these links ports can be clocked more than once per instruction cycle to achieve higher data transfer rates [DSPA94].

The Texas Instruments TMS320C80 DSP μ P is actually 5 separate processors integrated together on a single chip. The 5 processors consist of a single master processor and 4 slave parallel processors. The TMS320C80 Master Processor (MP) is a 32-bit, IEEE-754 compatible floating point processor. The MP manages the functions of all components within the chip. It is the main supervisor and distributor of tasks within the chip and also functions to communicate with external processors and service external interrupts. The 4 other Parallel Processors (PP) are 32-bit fixed point integer units. The 4 processors function independently of each other, perform data computations and handle operations associated with image, graphics and audio processing.

The C80 features 50 Kbytes of on-chip RAM and crossbar switching which allows the on-chip memory to be accessed by the PPs, MP and the Transfer Controller (TC). The Transfer Controller is an interface between internal and external memories and is used for cache servicing and data block transfers between external memory and internal SRAM. The C80 also features a Video controller that acts as a dual frame interface between the TMS320C80 and an image and display system. Finally, the C80 architecture includes a JTAG Emulation Interface that lets the user test programs or scan results through in-circuit emulation [TIMA94].

The Motorola DSP96002 was the first 32-bit microprocessor to implement in hardware the IEEE 754-1985 standard for binary floating point arithmetic. The DSP96002 includes 1 KW (kiloword) of program RAM, 1 KW of data RAM, 1 KW data ROM, a dual address generation unit, and a dual-channel DMA (direct memory access) controller. Two complete sets of 32-bit buses or ports are provided for external interfacing. Including this interfacing, the program and two address spaces can be expanded to 4 GW each with each comprised of 32-bits [GEOR92]. The DSP96002 also provides a Host MPU/DMA Interface for each of its external bus interface ports. Each Host Interface can be configured as a 8, 16, 24, or 32-bit wide parallel port which may be connected directly to the data bus of a host processor [MOTO89].

The following table summarizes some of the key features of the DSP μ Ps under consideration.

	TMS320C80	DSP96002	TMS320C80	ADSP21060
DIVIDE	45ns	233ns	125ns	150ns
PRECISION	32bits	32bits	32bits	32bits
RAM/ROM	1 Mbit	8Kbytes	50Kbytes	4Mbytes
1 CYCLE TIME	5ns	33.3ns	25ns	25ns

It should be noted that several of the microprocessors actually represent families of processors and therefore have the option of varying processor speeds. In such cases we consider the processor with the highest clock rate and therefore compare the highest peak performances provided in the literature for each processor. Digital signal processors are designed to perform the multiply and accumulate operations very efficiently since these operations are most prevalent in DSP applications. The processors under discussion take one instruction cycle to perform a single precision multiply or add. Furthermore, it was assumed that all the processors possessed the usual pipe-lined architecture required to perform a multiply and add (or multiply-accumulate) in one instruction cycle

3. Complex FFT Evaluations

The evaluation of DFTs (Discrete Fourier Transforms) with complex FFTs (Fast Fourier Transforms) is such an important component of DSP applications that most DSP μ P manufacturers include FFT performance data in the processor literature. The FFT Performance Graph of Figure 1 was constructed for the DSP μ Ps in terms of their complex FFT evaluation times for a 1024 point FFT. Where required, manufacturer's data values were interpolated

using the principle that FFT processing time is proportional to $N \log_2 N$ where N is defined as the number of points to be processed.

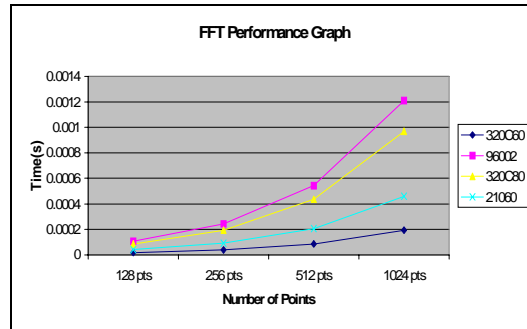


Figure 1 – FFT Evaluation Times for Selected DSPμPs

The significance of the performance graph in Figure 1 is best understood in terms of some simple illustrations. For example, the number of points in an FFT evaluation will determine the evaluation time for a specific DSPμP. For real time applications, the FFT evaluation time must not exceed the allotted time for transforming data within each DSP system algorithm sample time; i.e., the time to complete one cycle of all operations for a proposed real-time DSP system.

Thus, an analyst must be able to partition all DSP system functions and complete the evaluations within the allotted times - otherwise the proposed real-time system design would not be feasible; and clearly, FFT evaluation times are of prime concern for such applications.

4. IIR Filter Implementations

Now consider evaluating the performance of the processors in terms of IIR filter implementations. The times required to evaluate biquad IIR filter sections were 40, 100, 200, and 233 ns (nanoseconds) respectively for the C60, C80, 21060, and 96002. Neglecting overhead, higher order filter implementation times may be interpolated from the given data by observing that they are linear functions of the biquad implementation times.

Using the manufacturers DSPμP IIR filter implementation times, maximum sample rates for each of the processors may be calculated. The calculated sample rates for IIR filters up to 10th order are presented in graphical form in Figure 2.

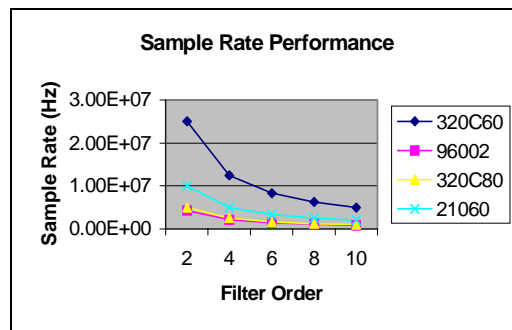


Figure 2 – Maximum Sample Rates vs. Digital Filter Order for selected DSPμPs

To interpret Figure 2, consider the requirements of an IIR 10th order filter design and implementation. The performance curves immediately indicate a maximum sample rate achievable by a specific DSPμP. If the filter bandwidth-to-sample frequency implies the DSPμP cannot meet the processor speed requirements, then another processor or multiple processors must be considered as alternatives to meet proposed design specifications.

As mentioned earlier, these sample rates, processor speeds, and filter complexity trade-offs are very similar to the issues of designing any DSP system. Thus, the information contained in Figure 2 implies key design issues in DSP applications other than just filter design.

5. Image Processing Benchmarks

A couple of important remarks should be made before the image processing techniques are presented. There have been a number of unsuccessful attempts to produce standardized functions for image processing benchmarking. Therefore if benchmarking is to be a factor in processor choice, it is up to the analyst to devise appropriate applications for processor comparison. Exact solutions to benchmark completion times can then be obtained by code implementation on each of the processors under comparison. Even with the aforementioned approach, execution times will still be heavily dependent upon the user's implementation and the software tools used to convert high level languages into code suitable for each processor. Taking these facts into account, the figures and times included in this paper are not intended to represent absolute values and will differ between implementations. However, the times presented were calculated in a manner consistent with reducing compiler and implementation dependency as much as possible.

Since obtaining exact solutions can become quite expensive, some benchmark completion times were interpolated based on existing comparisons and processor literature. Using this method provides a reasonable approximation to the completion times without introducing excessive cost and in some cases relieves some of the compiler dependency.

All of the applications chosen operate on 2-dimensional image matrices, with each element of the matrices representing a pixel value. The image matrices usually require more on-chip memory than is available and therefore external storage is implied in these cases. Since external memory accesses require additional instruction cycles, it is usually beneficial to limit these types of accesses as much as possible. The on-chip memory capacity of the ADSP21060 makes it capable of complete on-chip image and program storage in most cases. Thus, when compared with the other processors under consideration, the instruction cycles needed for data accesses by the ADSP21060 should be considerably less. However, for purposes of simplicity, in this treatment it has been assumed that internal and external memory accesses require an equal number of instruction cycles.

In addition to these arguments for formulating a reasonable approach to benchmarking DSP image processing algorithms, we include the two following assumptions.

1. The overhead time associated with partitioning an application into parallel jobs can be neglected when compared with the total application execution time.
2. We consider the instruction cycle time of the four parallel processors of the TMS320C80 to be equal to that of the master processor.

Using these two stated assumptions, we will consider the parallel implementation on the C80 whenever possible in an effort to give credit to this processor as an image processor since it was, to some extent, specifically designed for image processing algorithm implementation.

2-D FFT Evaluation

The basis of most frequency domain techniques in image processing begins with the application of the two dimensional Fast Fourier Transform. The information obtained by the application of the 2-D FFT can be interpreted to aid in edge detection and enhancement, brightness equalization using homomorphic filtering, and in some instances image compression.

The separability of the transform implies that the result may be obtained from successive calculations of 1 dimensional transforms and therefore partitioning can significantly reduce the algorithm complexity. It also implies that we may interpret 2-D FFT completion times from 1-D FFT data. Therefore, using the times given in the section on 1-D FFT evaluation, we obtain the graphical data depicted in Figure 3.

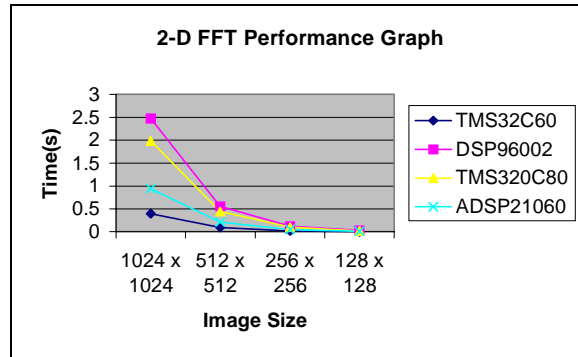


Figure 3 – 2 Dimensional FFT Evaluation Times

Brightness Equalization

Brightness equalization using histogram methods will be used to illustrate image enhancements using point-processing techniques. These techniques are based solely on the gray level value of a single pixel. Rather than specify a target histogram, we choose to implement the method that results in an image with a uniform histogram. Thus, the resulting image should have a greater dynamic range of pixel values and exhibit a higher degree of contrast. All images under consideration are represented using 256 gray levels.

The histogram equalization is comprised of the 3 following parts:

1. Calculation of the intensity percentages in the original image
2. Calculation of the pixel value transformations based on the original intensity percentages
3. Pixel value replacement based on the calculated transformation.

The numbers of cycles needed to perform a conditional add and a data comparison and replacement can be determined from the user manual of each individual processor. Using these values, the processing times for histogram equalization were calculated and used to construct the Histogram Performance Curves in Figure 4.




Figure 4 – Histogram Completion Times vs. Image Size

Spatial Processing

Finally, we illustrate spatial processing through the use of calculation of the Sobel Operators. Spatial image processing implies that the pixel intensity value after processing is determined by the values of the neighboring pixels. Calculation of the Sobel operators allows us to perform edge detection and enhancement by approximating derivatives. Separately calculated, the derivatives in each direction can be approximated by two 3x3 masks and provide the vertical and horizontal edges of an image [GONZ92]. The images may then be added to provide the edge-detected image.

Application of the operators simply consists of 6 multiplies and 5 adds. Since the masks must be applied to every pixel in the image, the total number of instruction cycles needed to determine the vertical or horizontal edge component is 6 instruction cycles times the number of pixels in the image. The number of pixels in the images will determine the number of additional add cycles needed to calculate the final image. Neglecting thresholding, the processing times can be calculated and depicted graphically as in Figure 5.

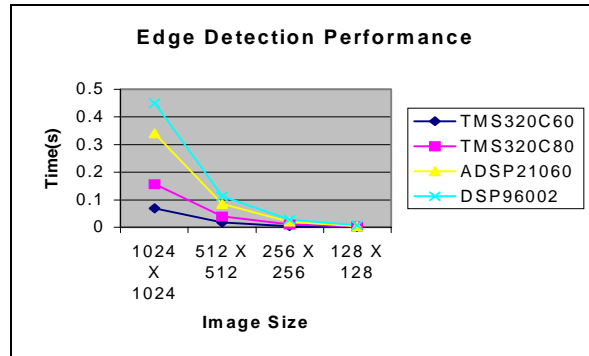


Figure 5 – Edge Detection Performance Using Sobel Operators

Thus, the image processing capabilities of the four DSP μ P candidates can be measured to some extent in terms of their ability to perform the basic frequency, point, and spatial applications included in this treatment. In view of the assumptions and the diversity of image processing, a variety of graphical benchmarks would result – hopefully without excessive contradictions.

4. Conclusions

The wide variety of features associated with the processors chosen allows us to illustrate the importance of the application in processor choice. Even though manufacturers continue to increase the speed and number of processor elements per DSP μ P in an attempt to make them more general purpose, no digital signal processor is best suited for all applications. However, we can draw some general and specific conclusions concerning the benchmark evaluations of the candidate DSP μ Ps.

1. The speed of the fixed point TMS320C60 represents the fastest processor according to our evaluations and therefore it should be given consideration for many applications. Conversely, its applicability is limited by the fact that it is a fixed-point processor and therefore processor accuracy becomes a factor for applications with a requirement to process variables with wide dynamic ranges. However, the processor is well suited for the most of the image processing applications presented in this discussion.
2. If accuracy as well as speed is an issue in processor choice, the ADSP21060 should be given consideration in view of its performance evaluation. Since this is the case in many DSP applications, the processor should be given serious consideration for a wide range of applications. Also, given its large on chip memory capacity, the ADSP21060 should be competitive with the C60 and C80 families in image processing applications.
3. Most of the features of the TMS320C80 have been designed for video and image processing applications. The system architecture also allows for ease of interface to other video processing equipment. Since the C80 already incorporates interfaced processors, it should also be included in applications were parallel implementation is a possibility.
4. The DSP96002 is also a good candidate for parallel application implementation. Linear arrays of processors are easily implemented with the DSP96002 and the “hooks” for testing and implementing these arrays are included in the software tools. In fact, the designers of the DSP96002 provided the superior environment for multiprocessor architectures. Even though the processor has the slowest speed of the processors included in the discussion, its special 96-Bit precision make it the processor of choice if accuracy is the prime issue.

5. It is not surprising that the performance evaluations generally tend to coincide with the chronological development of the candidate DSP μ P's; e.g., the TMS320C60 is the newest of the DSP μ P's with the DSP96002 being the oldest.
6. The proper choice of a DSP μ P for a specific application not only depends on the processor speed. Other issues concerning the processors architectural adaptability to the specific applications can also be critical.
7. Finally, if resources permit, seeking an optimum algorithm for mapping DSP algorithms to specific processors may represent the only alternative for the successful implementation of new complex applications.

7. Bibliography

[SIMO93] Simons, F. O. Jr.; George, A. D.; Short, C. W.; and Tait, H. S.; "*Architectural Tradeoffs in Digital Dynamic Model Implementations*", Proceedings of the TMS320 Educators Conference, Houston, Texas, August 10-12, 1993.

[GEOR92] George, A. D., *Microprocessor-Based Parallel Architecture for Reliable Digital Signal Processing*, CRC Press, Boca Raton, Florida, 1992.

[GONZ92] Gonzales, R. C., Woods, R. E., *Digital Image Processing*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1992.

[DSPA94] *DSPatch, Number 29*, Applications Newsletter, Analog Devices Inc., Norwood, Massachusetts, 1994

[TI6097] *TMS320C62xx Technical Brief*, Texas Instruments, Houston, Texas, 1997.

[TIPR97] *TMS320C62xx Product Review*, Texas Instruments, Houston, Texas, 1997.

[MOTO89] *DSP69002 Users Manual*, Motorola, Phoenix, Arizona, 1989.

[TI8095] *TMS320C80 Users Manual*, Texas Instruments, Houston, Texas, 1997.

[TIMA94] Uncopyrighted Work Presented by Texas Instruments and Marshall in the form of a seminar, 1994

[SIMO97] Simons, F. O., Jr. and Robinson, A. L., "*Performance Comparisons of Modern Digital Signal Processing VLSI Microprocessors*", Proceedings of the 105th Annual Conference and Exposition, Milwaukee, WI, June 15-18, 1997.

Keywords

Digital Signal Processing, Microprocessors, DSP Microprocessors, and DSP Benchmarks.

Biographical Information

AARON ROBINSON

Aaron Robinson received his BSEE and MS degree in Electrical Engineering from the Florida State University in 1994 and 1998 respectively. He has served as a project leader for the DSP group in the FAMU-FSU College of Engineering High Performance Computing and Simulations Research Laboratory. He is a McKnight Fellow committed to finishing his Ph.D. degree and becoming a professor of Electrical Engineering.

DR. FRED O.SIMONS JR.

Fred O. Simons Jr. received his BSEE degree in August, 1960 from Mississippi State University, and MS and Ph.D. degrees in electrical engineering from the University of Florida in 1962 and 1965, respectively. From 1965 to 1972 he was with the University of Florida. From September 1972 to September 1977 he was at the University of Central Florida. In 1978, after a years leave of absence to direct the Navy Lab, he joined the FAMU/FSU College of Engineering as a Professor of Electrical Engineering and as the Director of FEEDS(Florida Engineering Education Delivery System). He also serves as the Director of the High-performance Computing and Simulation (HCS) Laboratory at Florida A&M University and Florida State University. In 1987-88 he was the Interim Chair at FAMU/FSU. He is a member of Tau Beta Pi, Phi Kappa Phi, Eta Kappa Nu, IEEE, SCS , FES and ASEE.