

## *The Impact of Microprocessors on Medical Imaging Technology*

*Imaging systems employed in demanding medical applications such as computed tomographic reconstruction (CT) and multi-dimensional ultrasound typically require real-time or near real-time high-performance computing resources. While these systems have traditionally relied on proprietary architectures and custom components, recent advances in high-performance microprocessor technology have produced an abundance of low-cost components suitable for use in these high-performance computing systems.*

*A common pitfall in the design of high-performance imaging systems—particularly systems employing scalable multiprocessor architectures—is the failure to balance computational and memory bandwidth with I/O. The recent introduction of microprocessors with large internal caches and high-performance external memory interfaces makes it practical to design high-performance image reconstruction systems with balanced computational and memory bandwidth. In these systems, both the board level memory and I/O architecture as well as the microprocessor memory have significant performance impact. Systems that do not scale the memory bus bandwidth as processors are added do not improve in performance, or they reach a limiting value after some initial small gains in performance. In addition, the I/O bandwidth plays a significant role in overall performance.*

*Therefore, the most important factors in the selection of a programmable imaging system are: (a) will the system be “native” or will it use a co-processor model, (b) which processor has the correct balance of processing and I/O for the application, and finally (c) what is the overall cost? Two real-world examples of such designs will be presented, and their performance and cost will be analyzed.*

### **Description of Applications**

In order to characterize the advantages of the various state-of-the-art microprocessors, we will present two typical demanding medical image reconstruction applications. Included are a balanced I/O and computation application such as CT scan reconstruction, and a demanding I/O and computation application such as ultrasound. We will calculate the performance of the latest microprocessors from Analog Devices, Intel, Philips Semiconductors, and Texas Instruments<sup>1</sup>.

<sup>1</sup> The Motorola PPC750 AV (AltiVec) and the Hitachi/Equator MAP1000 are no

The important properties of the selected processors are

reach bus saturation very quickly to the point where adding additional processors will not improve the rate at which calculations can be performed. The table assumes that the processors are isolated from each other when they are

most multiprocessor co-memory designs due to scalability issues. The only non local memory design is the “native” PIII-comparison purposes.

SPECIFICATION	Intel PIII-450	Philips TM1300	TI C6701	ADI 21160
Architecture	CISC	VLW	VLW	VLW
FPU	Yes	Yes	Yes	yes
MFLOPs (Peak)	250	720	1000	600
MIPS (Peak)	450	900	1336	100
MOPS (Peak)	900	4860	1336	800
Memory Bus Bandwidth (MB/s)	800	572	332	400
1K FP cfft (μsec)	300	106	108	90
1K 16 bit cfft (μsec)	191	63	108	90
1K FP dot product (μsec)	7.38	2.84	3.07	5.12
16x16 MACs (MMAC/s)	2.32	1.42	3.07	5.32
8x8 MACs (MMAC/s)	9.34	6.55	7.11	11.80
512 <sup>2</sup> x8 bit Conv3x3 (msec)	5.34	2.62	7.11	11.80
512 <sup>3</sup> x8 bit Erosion/Dilation (msec)	7.34	1.42	3.62	3.93
“Glue” Logic Cost (\$/CPU)	\$150	\$3	\$65	\$39
CPU Price (\$)	\$300	\$100	\$150	\$100

### **Reconstruction Application**

The reconstruction region is a square array of 1024 x 1024 pixels centered on the axis of rotation of the x-ray head (Figure 7). Data is collected from the x-ray head as radial scans, with 4 channels of 4096 values for each of 360 angles. The data collected occupies 4\*4096\*360\*2 = 11,796,480 bytes. The data is moved into memory at an average rate of 5 megabytes per second (Figure 8).

<sup>2</sup> & <sup>3</sup> Herman, GT. “Image reconstruction for projections”, Academic Press. 1980.

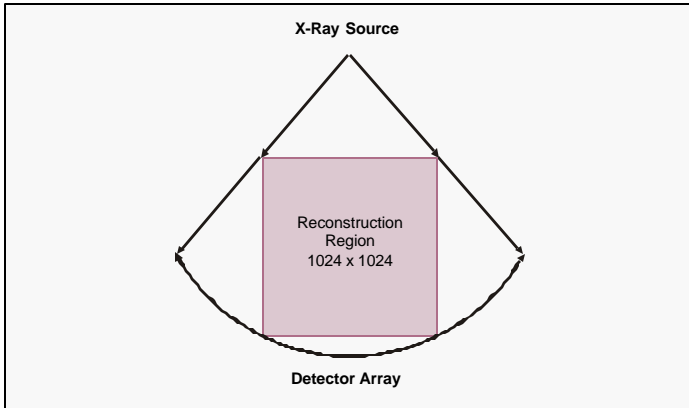


Figure 7. Relation of the region of reconstruction to the x-ray source and detector array.

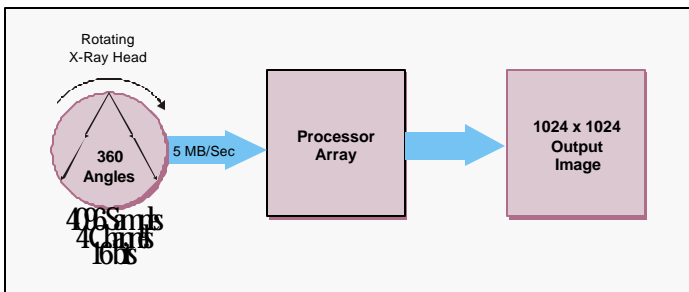


Figure 8. Data collection.

The following steps reconstruct the pixel densities:

1. Calculation of the density estimate from the output of the four detectors at the end of each beam path.
2. Back-projection. Conversion of the fan-shaped polar form of the data to the square array of pixels in the reconstruction region is accomplished by summing the density along each ray into each pixel the ray encounters. Linear interpolation is used when the ray passes between two pixels (Figure 9).

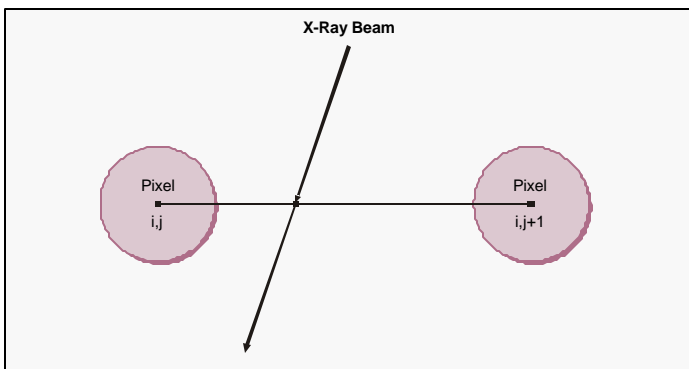


Figure 9. Back Projection.

Four channels of data are used to estimate the actual density along the beam path to compensate for the normal beam hardening that occurs as x-rays pass through an object.

The value summed into each pixel is:

$$p[i, j]+ = (1 - \alpha) D_{beam}$$

$$p[i, j + 1]+ = \alpha D_{beam}$$

$$\alpha = \frac{x_{pixel[i, j+1]} - x_{beam}}{\Delta x}$$

where  $D_{beam}$  is the density measured for the current beam being processed and alpha is determined from the geometry of the beam.

3. Application of a 2D FFT to the image array.
4. Application of a weighting factor of the distance of each pixel from the axis of rotation. This compensates for the greater number of rays passing through pixels closer to the axis of rotation.
5. Application of an inverse 2D FFT to the whole array.

Steps 2 through 5 effectively apply a convolution to the data to correct for the weighting error generated by the beam geometry, giving a good approximation to the density distribution of the object positioned in the reconstruction area <sup>2</sup>.

**Algorithmic Complexity**

The number of operations performed in each step can be counted as follows. These are summarized in Table 2.

Step	MACs	Adds	Reads	Writes
1	29,491,200	4,423,680	41,287,680	1,474,560
2	1,835,240,760		3,670,481,520	1,835,240,760
3	28,800,000		2,097,152	2,097,152
4	1,048,576		2,097,152	1,048,576
5	28,800,000		2,097,152	2,097,152
Total	1,923,380,536	4,423,680	3,718,060,656	1,841,958,200

\*Based on averaged 230 microsecond FFT time.

Table 2. Operations required

1. **Density Estimation from Detector Data:** A fifth order polynomial is applied to the output of each the four detectors at the end of each beam path, and the results of these four polynomials are summed to give the density estimate. This step requires 4\*(5 Multiples + 5 Adds) plus three adds per beam path, or 81,920 MACs (multiply/accumulates) and 12,288 adds.

For 360 angles, the total number of MACs is 29,491,200, and the total number of adds is 4,423,680. Additionally, each detector has an array of 24 coefficients and 4 values that must be read from memory, followed by a write of the

density estimate. This results in 41,287,680 reads from memory, and 1,474,560 writes to memory to process all 360 angles.

2. **Back-projection:** For each angle, one third of the beam path passes into the top and out from the bottom of the reconstruction region, encountering 2 pixels for each of the 1024 lines it crosses. The weights for each beam path is pre-computed, and so each of the 2048 pixels on these 1366 beams requires one MAC, or a total of 2,797,568 MACs. The other two thirds of the beam paths access from 3 to 2048 pixels depending on how far up the left or right side of the reconstruction region the ray path lies. In total, these beam paths access 2,300,323 pixels, requiring 2,300,323 MACs. For all beams, 5,097,891 MACs are needed. Each MAC requires a pixel value to be read from memory, a table value to be read from memory, and a final write of the new pixel value back to memory. For all 360 angles, the back-projection requires 1,835,240,760 MACs, 3,670,481,520 reads from memory, and 1,835,240,760 writes to memory.
3. **2D FFT:** A 1024 x 1024 2D real FFT is composed of 2048 real 1024-point FFTs, 1024 along the rows and 1024 along the columns. Two real signals can be processed concurrently with one complex 1024-point FFT plus 2,044 additional adds. Two real FFTs can be performed by this method in 460 microseconds on a single 40 MHz SHARC 2106x processor, averaging to 230 microseconds for a 1024 point real FFT. The whole 2D FFT thus requires 0.470 seconds plus, plus 0.013 seconds required to read and write the whole array twice from memory, or 0.483 seconds on a single processor.
4. **Pixel Weighting by Radial Position:** A pre-computed table of radius values is read with the image data, multiplied and written back to memory. This step requires 2,097,152 reads from memory, 1,048,576 multiplies, and 1,048,576 writes to memory.
5. **Inverse 2D FFT:** The 2D inverse FFT requires the same number of operations as required by the 2D FFT in step 3.

**Parallelization of the Algorithm**

In order to improve performance, the computation is divided between processors by splitting the reconstruction region into horizontal stripes. The first processor will get the first N rows (N = 1024 / number of processors), and each processor will get succeeding rows. The algorithm is inherently linear, allowing this natural division between processors.

**Memory Requirements**

This algorithm uses several pre-computed tables. Each detector has a table of 24 calibration coefficients for a total of 98,304 four-byte floating point numbers. The pixel radius value table (distance from the rotation axis) requires 131,072 floats, taking

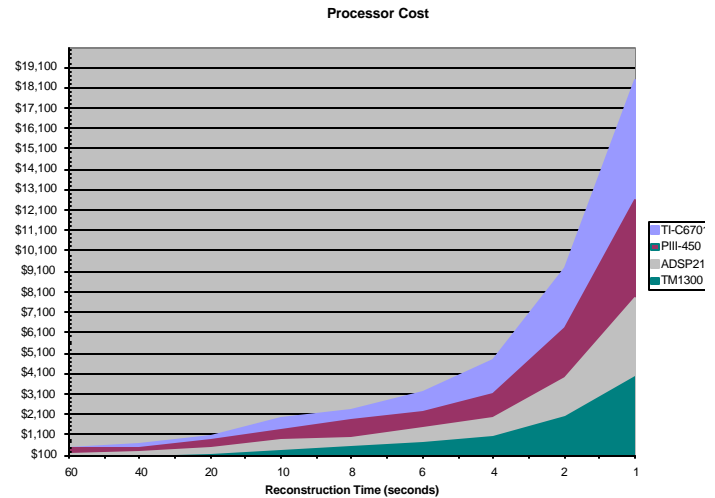
advantage of symmetries. Again taking advantage of symmetries the ray table requires 57,351,273 floats. The reconstructed image requires 1,048,576 floats, and the input data requires 11,796,480 bytes. The total memory requirement is therefore under 256 megabytes.

Total (Elapsed) is the time the routine takes to run on one processor with overlapping I/O and compute.

Processor	Total	Time (Mac)	Time (Add)	Time (read)	Time (write)
TM1300-180	38.88	2.67	0.006	26.00	12.88
TI-C6701	85.18	5.79	0.013	57.27	27.91
ADSP21160	55.60	9.62	0.022	37.18	18.42
PIII-450-MMX	27.78	8.55	0.009	18.58	9.20

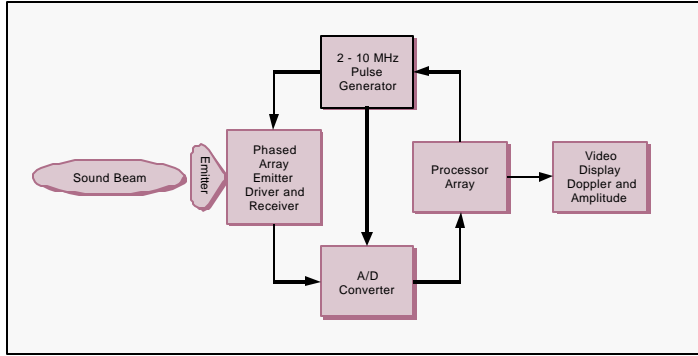
**Processor Cost vs. Reconstruction Time**

The illustration on the next page shows the cost/performance analysis for the four processors.



**Ultrasound Application**

Processing power available from many DSP processors has allowed a reduction of the amount of analog processing required in Ultrasound applications. In addition, they have reached performance levels that will allow additional features to be implemented without a significant increase in cost.



Block Diagram of Ultrasound System

Visual and auditory feedback is critical to the placement of the probe, which is extremely important to the quality of the test results. To assist in this difficult task, the ultrasound system must process the data it collects and render an image as quickly as possible. This would allow the technician to receive essential visual feedback in the placement of the ultrasound probe without significant lag. This requirement generates a tight latency requirement on the processing of ultrasound data.

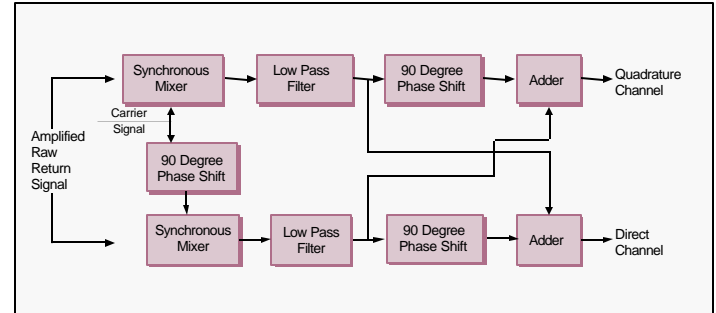
A typical ultrasound machine might take data at 20 million 12-bit samples per second on two channels. This rate allows operation at carrier frequencies of 10 MHz or less, which is typical for ultrasound systems.

Several processing steps must be performed on this high sample rate data to reliably detect the signals generated by the sound returns. Two components of the signal are of interest: the **amplitude of the reflection** and its **frequency**. The amplitude of the return is used to detect density of tissue, while the frequency is used to detect motion via the Doppler effect.

The emitted sound beam is generated from a phased array emitter, which is pulsed with a carrier frequency of from 2 to 10 megahertz. The phased array emitter allows the sound beam to be positioned electrically ('beam steering') and focused into a small area ('beam forming'). The sound reflected back is received by the emitter and passed to an A-D converter for conversion to digital form. The A-D converter processes the return signal and the original emitted signal generating two streams of digital data—a signal and a quadrature channel.

The digital data is processed in two ways. First, the amplitude of the signal return is demodulated from the high frequency carrier with a synchronous AM detector. Second, the phase of the returns is compared to the originally emitted carrier, with a synchronous quadrature detector. This phase measurement is combined with other phase measurements to detect the Doppler shift in the return signal. In addition to these demodulation steps, the signals are gated by time and angle to report data from a region selected by the operator. Additional processing can be performed to remove artifacts such as echo returns,

variable sound velocities, and geometrical artifacts (e.g. off angle Doppler measurements).



The quadrature phase detector is implemented in software on the processor array. The pair of synchronous mixers and the low-pass-filters operate at the 20MHz-sample rate. The output of the low-pass-filters is sample rate converted down to a 25 kHz for further processing (see diagram above).

To implement the high sample rate processing, the following steps are required: a multiply in each synchronous mixer; a delay for the 90 degree phase shift; and a multistage rate-converting low pass filter. The filter is implemented as a three tap filter, followed by a rate conversion down to 1MHz, then another three tap filter and a rate conversion down to 100kHz, then a 17 tap filter and a rate conversion down to 25 kHz. The rate conversions reduce the computational load and allow a sharp cutoff filter implementation. A non-rate-converting filter would require significantly greater processing bandwidth. Each filter tap involves a multiply and an add. The phase shifts are essentially free as they are implemented through modifying the address from which the data is taken. Totalling the computational requirement, we have 8 multiplies and 6 adds at 20 MHz, then 6 multiplies and 6 adds at 1 MHz, then 34 multiplies and 34 adds at 100 kHz. The required performance is  $(280+12+6.8=298.8)$ , about 300 MFLOPs. This is well within the reach of DSP processors on the market today.

After this processing, the two channels are combined to form the Doppler shift signal, which is gated and then processed with an FFT for display. The gating selects a window determined by the operator, in which signals are taken and out of which signals are ignored. This involves varying the data selection window, which translates into a start and stop time, for each beam position.


The demodulated amplitude signal is corrected by applying an increasing gain with time to correct for attenuation, and warped to correct for speed variations. The amplitude signal is plotted on the display at the locations dictated by the beam position. This plot creates the image of the area being studied.

These post processing steps, except for the FFTs, are insignificant when compared to the processing needed by the demodulator and the FFTs. The FFT processing is performed to

display the spectrum of the Doppler return, which translates directly into velocity, after considerations for geometry. The character of the spectrum (noisy, smooth, wide band, narrow band) is an indication of the condition of the area being examined.


The processing required by the FFTs depends on the rates selected by the operator. The 25 kHz Doppler signal can be rate converted up to give finer frequency velocity resolution, and can be performed lapped (i.e. overlapping sliding buffers) to improve low velocity sensitivity of the system. All of these effects translate into an equivalent number of 1K complex FFTs per second. The table and graph below show how well various processors perform the FFT processing.

Two considerations are significant to this processing—the rate at which data is passed to the process performing the FFTs and the bus architecture used. As the number of FFTs per second increase, the computational demand on the processor increases. At the same time, the demands of the processor for data increase as well.



This results in the following requirements for processors assuming almost linear scaling. This is true for most designs except the “native” PIII-450 case. Thus DSP processors available today are now fast enough to assume responsibility for the signal processing performed in ultrasound systems at the carrier frequency, eliminating the requirement for special purpose demodulators implemented in hardware.

FP 1024 CFFT/sec	TM1300	ADSP2116X	TMS320C6701	PIII-450
5000	2	1	2	2
7000	2	2	3	3
10000	3	2	4	NP
20000	5	4	8	NP
50000	11	9	20	NP
70000	15	12	28	NP
100000	21	17	40	NP



Notably, an excellent DSP processor, the Intel PIII, is limited by its surrounding logic (the PC) and is unable to perform in some applications. Although the use of the AGP bus would improve the situation, its SMP design will ultimately limit its scalability. Therefore the most practical solution for demanding application remains a co-processor board that is more scalable, has higher throughput, and ultimately is cheaper than the native solution. In this group the TM1300 from Philips has excellent performance vs. cost under most circumstances.

### Conclusions

Example applications have been presented illustrating how memory bandwidth and processor performance limits the performance of multiprocessor systems. Whereas memory bus saturation severely limits the scalability of cluster based architectures (e.g., PIII-450), the local memory architectures of co-processor boards allow throughput to scale linearly with the number of processors. The currently available processor are sufficiently powerful to perform demanding medical imaging functions while maintaining significant cost competitiveness with custom solutions.

\*Correspondence: Email: [joe@alacron.com](mailto:joe@alacron.com); <http://www.alacron.com>; Fax (603) 891-2750



71 Spit Brook Road, Suite 200  
 Nashua, NH 03060  
 Tel: 603.891.2750 Fax: 603.891.2745  
 Web: [www.alacron.com](http://www.alacron.com) E-mail: [sales@alacron.com](mailto:sales@alacron.com)

© Alacron, Inc. All Rights Reserved.