



35 South Service Road • Plainview, NY 11803
TEL: 516 694-6700 • FAX: 516 694-6715

The Future in Microelectronics

CIRCUIT TECHNOLOGY

APPLICATION NOTE #101

A Proposal for Low Cost X-Ray Tomography Using R4430 Microprocessor MCMs as Parallel Processing Elements

Aeroflex Laboratories Inc. teamed with the State University of N.Y. at Stony Brook

Technical and Administrative Points of Contact:

Aeroflex Laboratories Inc., 35 South Service Road, Plainview N.Y. 11803

Tom Terlizzi, TEL: (516) 752-2418, FAX: (516) 694-6715

State University of New York at Stony Brook

Larry Wittie, TEL: (516) 632-8456, (email: lw@sbc.suny-sb.edu).

A. Innovative Claims for the Proposed Research

The Aeroflex-Stony Brook team is proposing the development of a demonstration model parallel computing system with the capability of providing the computing power required for real time X-Ray tomography. Low cost RISC microprocessor multichip modules will be used as the basis of the design. The system design will utilize microprocessors linked by innovative latency-hiding cache-update interfaces for the fast eager copying of individual changes to variables shared between processors. The system allows rapid writer-initiated sending of cache lines to new processors that may not yet have a copy of the cache line. Compile-time analysis tools will be used to determine the minimum data sharing between processors sufficient to produce accurate results.

The proposed processor assembly will contain 256 microprocessor modules, 256 SRAM (8 megabytes) modules and associated control and interface modules. An 17" by 25" by 2" (43 cm x 63 cm x 5 cm) system can deliver 8 GFLOPS, 2/3 of its peak power of 12 GFLOPS, to parallel FFT computations, allowing real-time calculation and display of high resolution medical tomographic scans. A 256 processor chassis with internal convective cooling would market for less than \$200,000 in production quantities. It would equal the performance of a 16 processor Cray YMP for FFT applications.

B. Deliverables Associated with the Proposed Research

Deliverables will be demonstration model hardware and software for a commercially viable real-time, X-Ray tomography system, a parallel computer capable of rendering 512 X 512 X 512 voxels in 1/6 second or less, or 256 X 256 X 256 voxels in 1/24 second or less.

The processor demonstration model will be a prototype for scalable low cost parallel computers with up to 256 microprocessors. The software will demonstrate the effectiveness and efficiency of the "Eager sharing" concept.

C. Technical Rationale, Technical Approach

The recent commercial availability of high speed (50 to 100MHz, 100 to 200 MIPS, 33 to 66 MFLOPS) RISC microprocessors and advanced low cost MCM "D" technology, has made a series low cost and high speed multiprocessor systems feasible. The R4430MC RISC multiprocessor designed by MIPS is our current choice for fast processors. It operates with a 3.3 volts power supply for low power dissipation. Speeds are expected to exceed 1000 MIPS per processor within the next five years.


The basic concept proposed is a low cost microprocessor module along with associated memory that can be stacked vertically with heat sinking along two edges and electrical connections on the other two. Stacks can be tiled horizontally for parallel processor applications. The modules are 3" L X 2" W X 1/8" Z in size and are cooled by conduction to the cooling tower edges.

Each stack contains the following modules

1. Four 50 GFLOPS, 75 MHz, R4430 modules containing a microprocessor, one megabyte of secondary cache SRAM, assorted buffers and PLL circuitry on an aluminum-based MCM substrate
2. Four memory modules each containing 8 Mbyte of shared SRAM
3. A DMA module for direct memory access to local shared memory for input from sensor arrays, or output to external display systems.
4. Two 4 way link controller modules for rapid passing of single cache lines to neighboring stacks.

The modules in each stack are connected by a 25 MHz, 64 bit wide memory access(Z) bus capable of passing 160 Mbytes per second in 4 word (256 bit) cache lines. Each cache line has a 64 bit address, of which the upper 20 bits are used for routing of cache lines among stacks and the lower 32 bits are used for addressing cache lines and shared memory blocks anywhere in the multistack system. Each stack is connected to four pairs of 25 MHz, 64 bit wide interstack links, all of which form an end-around mesh of 8 by 8 intersections for 64 stacks and 256 processors. Figure 1 shows the proposed modules.


FIGURE 1. Multiprocessor MCM Stacks



The microprocessors are linked by innovative latency-hiding cache-update interfaces for the fast eager copying of changes to variables on individual cache lines shared between processors. A fast bus joins the caches within each vertical stack to interstack links. Four way mesh parallel connections link neighboring stacks for eagersharing and remote memory accesses. The planned custom ASICs will allow 1.7 microsecond or shorter delays for cache lines sent after a local write to be received by the most distant processors sharing the new values. Seitz' router chips and SCI gigabit interfaces are also being considered for the interstack links. Write orders are preserved across groups of processors to allow safe low overhead synchronization using availability flags distributed with shared data.

Fast Fourier Transforms (FFT) and their inverses (IFFT) are needed for the signal processing to integrate X-Ray projections from multiple angles to form a 3-dimensional tomographic model. Preliminary simulations indicate that for FFTs, the fast cache-update (eagersharing) interfaces will allow nearly linear speedup (at roughly 2/3 of peak processor rates) for ensembles from 1 to 2048 computer modules. This is in sharp contrast to more traditional demand fetch mechanisms, which do not deliver improved performance in ensembles of more than 16 processors. For a 256x256 FFT calculation, eagersharing among 64 processors is 7 times faster than demand fetch; for 2048 processors, eagersharing gives results 160 times faster. Figure 2 compares speedups for eagersharing and demand fetch for 1 to 2048 computers for 256 x 256 FFT calculations

FIGURE 2. Network Power for FFT with 65536 Data Points



Other simulations show that the MCM semiconductor junction to heat sink temperature rise can be held below 33°C with an aluminum type “D” substrate and internal convective cooling. This substrate design provides a low cost assembly technology for commercial applications. Processor (hot) and memory modules (cold) are alternated to spread the heat more uniformly. The use of 3.3 Volt power supply die and fast SRAM devices limits maximum power dissipation per microprocessor module to 7 Watts maximum. This is much less than previously designed R4000 type MCMs. The average power dissipated per module will be 4.5 watts. Even with this reduced dissipation, 1200 watts will be generated in the chassis under peak conditions and efficient heat removal will be required. It is anticipated that standard commercial liquid cooled cold plates from companies such as R-Theta of Buffalo N.Y. or McLean Engineering of N.J. will be adequate for the task. A simplified cross section of the heat transfer path is shown in figure 3.

FIGURE 3. Module Cross section





Figure 4, shows a top view of the chassis and modules

FIGURE 4. Modules and Chassis, Top View



D. Conclusion

The development of the 256 parallel processor demonstration model will stimulate market insertion of multichip modules in the following ways;

- This processor will be capable of performing real time, high resolution tomography in a small physical size, at a reduced cost. It can be used as a prototype for the development of a commercial piece of hardware for medical, industrial and scientific applications.
- The software and hardware developed for this application will be scalable up to 256 processors and will provide stimulation for the application of multichip modules to parallel processing problems.
- Aeroflex will offer the Processor MCM, Memory MCM and associated ASIC devices designed in this program as standard product building blocks for applications in commercial and industrial applications. It is anticipated that the selling prices for the Processor MCM will be in the order of \$500 in production. The memory module and ASICs will also be attractively priced to stimulate industrial and scientific application.
- Since Aeroflex has considerable experience in selling MCMs to the military market over the years with Standard Military Drawings and custom circuits, it is anticipated that these MCMs can be converted to military qualified designs and be offered as standard products for military systems.

- The use of the R4430 RISC processor and high speed SRAMs as die for this program will stimulate the application and development of "Known Good Die" for these devices and accelerate the application of RISC microprocessors and SRAM die for military and industrial MCMs. See Figure 5
- The use of multiple sources (currently six semiconductor vendors) for the R4430 MC die will drive the price to commercial cost levels for the proposed parallel processor.
- The use of the low cost (\$10/ sq in) thin-film-polyimide-aluminium substrate in a production mode will gain valuable experience in the application of lower cost type "D" substrate technology for industrial applications.
- Aeroflex will utilize the Hewlett Packard HP 82000 in the testing phase of the modules. This high speed tester and the software methodology can be scaled to 200MHz clock rates for advanced RISC processor and memory modules. This will accelerate the "Known Good Die" program for MCM products.




FIGURE 5. Cofired Ceramic R4400 with 256K Secondary Cache

E. Comparison with Related Research in This Area

The Aeroflex-Stony Brook research is a substantial improvement on current research in this area. For example, the DARPA contract, No. J-FBI-91-280 placed with the Microelectronics and Computer Technology Corporation (MCC) is for the development of RISC processor modules for commercially scalable parallel computing systems. A paper presented at the 1993 International Conference on MCMs in Denver, April 14-16, 1993, by Robert Smith II and Paul Hunter, entitled "MCMs for Parallel Computers", (ICEMM Proceedings, 1993) describes the program and the current state of the research.

The Aeroflex-Stony Brook approach offers the following advantages:

- The use of 3.3 volt microprocessor and memory die results in a maximum power dissipation per module of 7 Watts and an average of 4 to 5 Watts. This is a factor three less than the dissipation of the modules proposed in the paper.
- The reduced power supply current and heat load of the Aeroflex-Stony Brook design allows for more compact design and reduced size and weight.

- The compact single chassis design reduces wiring delays and power dissipation due to driving excessive shunt cable and wiring capacitances.
- The modular approach permits ease of repair and trouble shooting during operation and during the development phase of the program.
- The use of the thin-film-polyimide-aluminium substrate technology (available from vendors such as MIC Technology an Aeroflex Division) provides a fast, low thermal resistance substrate for the micro-processor MCM with the resulting lower temperature rise and improved reliability.

Other systems on the market do not take advantage of the "Eagersharing" technology developed by Stony Brook University, with the result that the Aeroflex-Stony Brook system will utilize processors more effectively. Greater efficiency will reduce the proposed system cost, size and weight.