Defense Technical Information Center Compilation Part Notice

ADP010841

TITLE: An Embedded Fusion Processor

DISTRIBUTION: Approved for public release, distribution unlimited
Availability: Document partially illegible.

This paper is part of the following report:

TITLE: Space-Based Observation Technology

To order the complete compilation report, use: ADA391327

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, ect. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP010816 thru ADP010842

An Embedded Fusion Processor

John Rooks, AFRL/IFTC
26 Electronic Parkway
Rome, NY 13441-4514 USA

Summary:

This paper describes an embedded High Performance Computer (HPC) designed to perform the sensor data fusion for the Discriminating Interceptor Technology Program (DITP). The HPC's electrical and physical architecture will be reviewed. The processor's architecture, FPASP5, evolved from years of Ballistic Missile Defense Organization (BMDO), and US Air Force research into wafer scale packaging and power efficient programmable signal processors for space-based applications. The processors, memory, and interface bare chips are packaged in Multichip Modules (MCMs). Our current version is designated MCM3. These MCMs can be stacked in thin layers before being inserted into the chassis level interconnect scheme. The chassis interconnect leverages a BMDO and Air Force Research Laboratory (AFRL) funded technology called Highly Integrated Packaging and Processing (HIPP). HIPP allows MCMs and two by two inch Printed Circuit Boards (PCBs) to be stacked together and interconnected with printed flexible flaps and a micro backplane. The combination of these techniques allows us to meet the strict constraints of space-based surveillance and interceptor applications.

The following description starts with the processors followed by their interface chip and communication protocols. Then continues with the MCM and chassis level packaging. Finally the Fusion Processor's (FP) integration with other components and its software environment are reviewed.

Floating Point Application Specific Processor (FPASP5):

The FPASP5 was designed by AFRL's Information Directorate for compatibility with MCM packaging and for power efficient floating point operation. The processor was kept simple to allow radiation hardening of the design. It employs external commercial SRAMs for its primary memory within each computing element. Each die contains two processors that operate independently, for highest throughput, as a processor/coprocessor pair, for lowest latency, or as a self checking pair, for fault tolerance. Having a simple core processor keeps the die costs low and yields high. Each processor has a 64 bit static memory interface and shares with its die mate a 64-bit IOBus interface and a boundary scan test access port. All communication with the processor registers and memory from outside the die goes through the IOBus. Using the boundary scan test access port for bare die testing eliminates probing damage to the bond pads.

Each processor can perform two 32-bit multiplies and two Arithmetic Logic Unit (ALU) operations per clock cycle or one 64-bit multiply and one 64-bit ALU operation per clock cycle. The pair of processors on each die execute eight 32-bit operations per clock. Peak performance is 320 MFLOPS per die with the current generation. It is fabricated in a 0.5 µm CMOS process and runs at 40 MHz.

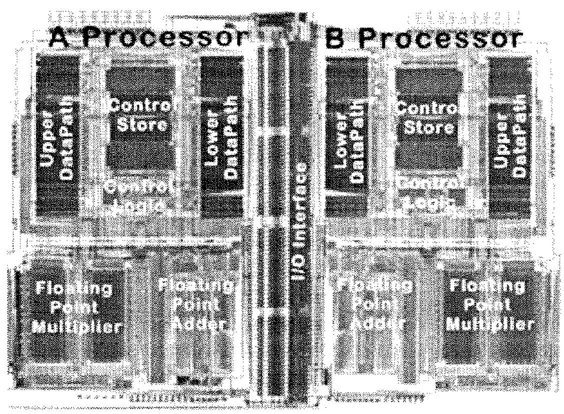

Figure 1: Floating Point Application Specific Processor

Each processor has an independent bus for its static memory. Additionally it is possible for each processor to read, write, or execute from its die mate's memory by temporarily shutting the other processor's clock off. The use of a relatively large Static Random Access Memory (SRAM) for each processor's main memory and the lack of a cache gives predictable execution times no matter where an operand is located in memory. And unlike Dynamic Random Access Memory (DRAM) there is no penalty for non-unit stride memory references frequently encountered in signal processing. Predictable memory performance simplifies program optimization. [1] [2]

In the MCM3 four processor chips (eight processors) are interconnected using their IOBus. The PCIF2.0 bridge chip among other things is the host for the processors on the IOBus, granting the bus and preempting processors to clear the bus when necessary.

Paper presented at the RTO SET Symposium on "Space-Based Observation Technology", held on the Island of Samos, Greece, 16-18 October 2000, and published in RTO MP-61.

PCIF2:

Messages that leave the MCM3 or originate from outside the MCM3 pass through the PCIF2 shown in figure 2. The PCIF2 in addition to being the IOBus host provides access to a 32/64-bit PCI bus, a Myrinet FI32 SAN interface chip, and Synchronous Dynamic Random Access Memory (SDRAM). Each interface has a non-blocking connection to each of the other interfaces, with all internal transfers taking place between First-In-First-Out memories (FIFOs) at the IOBus clock (or FPASP5 clock) rate.

There are eight FIFOs on the PCIF2. Four of the eight FIFOs receive data from outside the chip and the other four receive data from inside the chip (primarily other FIFOs). All FIFOs are eight bytes wide. The PCI and SDRAM FIFOs have asynchronous controllers allowing the IOBus clock to run independent of the PCI and SDRAM clocks. The interface to the FI32 chip uses the IOBus clock for convenience since the FI32 die provides the asynchronous interface to the Myrinet link. The FI32 version 1.3 has two independent channels one for inward bound messages and one for outward bound messages, each runs at 160 Mbytes per second

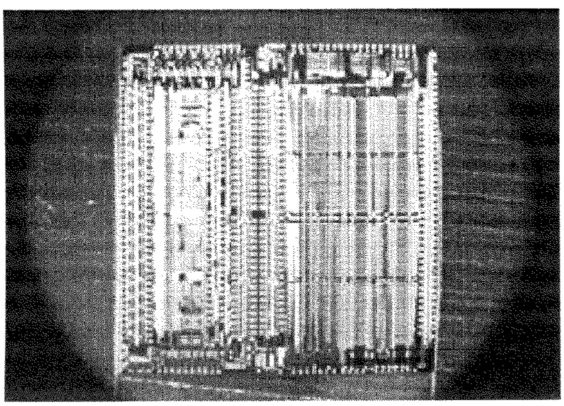

Figure 2: PCIF2.0

To maximize throughput the PCI outgoing FIFO is 8 bytes by 80 deep. Outgoing PCI messages can be stored and forwarded when the PCI bus is not available. The PCIF2 uses the PCI command "Memory Write and Invalidate" as well as the more standard "Memory Write". Standard Memory Writes cause the receiving processor to get an interrupt after the write is completed. The Memory Write and Invalidate command is received without interrupting the receiving processor. To prevent deadlocks and insure message progress is made, incoming PCI messages receive priority over outgoing PCI messages. Incoming PCI messages destined for the IOBus will preempt any current outgoing IOBus message to allow the incoming PCI message to get to its destination. The incoming PCI FIFO is 8-bytes wide by 16 deep.

Both Myrinet FIFOs are 8-bytes wide by 128 deep. This in combination with packets that are on the order of 1K-byte keeps the 160 Mbyte/sec Myrinet link from consuming excess IOBus time, since the IOBus transfer rate is 320 Mbytes/sec. Myrinet packets can exceed the 1K-byte. The only consequence is that after the FIFO is emptied the IOBus will be transferring data at 160 Mbytes per second, only using one-half of the IOBus's potential 320 MBytes/sec bandwidth. Using the 1K byte packet size allows the processing elements to keep both of the independent (in and out) Myrinet channels busy. To maximize data transfer a special packet was registered with Myricom for the FP and it associated sensors.

The AFRL Myrinet Packet (0B00 Hex) is used to reduce processing overhead and make Myrinet messages transparent to the application programmer. Various subtypes are defined. The primary subtype that the processors use is the WSSP direct write packet subtype (0000). This has the simple definition that when a packet of this type arrives it is forwarded by the PCIF2.0 to the address contained by the next 8-bytes, (those following the packet type and subtype). So an arriving packet of type and subtype 0B000000 can be written to any processor memory or register, the PCI bus, SDRAM, or the outgoing Myrinet port. If the message was addressed to a processor's SRAM it would first be held in the incoming Myrinet FIFO until the end of the packet is received or the FIFO is full. Then it would be forwarded to the outgoing IOBus FIFO where it would go onto the IOBus preempting any other activity on the IOBus except for an incoming PCI transaction. All messages except incoming Myrinet messages pass through the PCIF2 with only a few clock cycles of latency. The Myrinet messages are intentionally delayed until the end of the packet is received or the FIFO is full to maximize IOBus efficiency. [3]

The SDRAM is used to assemble smaller messages destine for the Myrinet and for data storage. For instance, a message traveling from the PCI bus to the Myrinet could become broken up as it traverses PCI-PCI bridges and other agents request the bus. By first dumping the data into the SDRAM and then issuing an instruction to the SDRAM controller to write the message to the Myrinet, you can insure that an intact Myrinet message is sent out. This is important since each Myrinet message requires the appropriate routing bytes prefixed to the packet type and packet data. When the SDRAM is used for data storage its data can be accessed by many processors without burdening an individual processor with multiple memory accesses. It is also capable of writing back its memory using a non unity stride which prevents unnecessary data transfers.

MCM3A:

General Electric's (GE's) High Density Interconnect (HDI) is used for our bare die interconnections. The GE HDI process uses copper interconnects separated with KAPTON (Dupont trademark) for a high density interconnection between the silicon chips and the package.

Our current version is designated MCM3A. It has four processor chips (eight processors), one PCIF2 interface chip, 16 8-Mbit SRAMs, 2 64-Mbit SDRAMs, a FI32 version 1.3 Myrinet chip, as well as capacitors and resistors. Figure 3 below shows that there is very little wasted space in the MCM3. The chips are interconnected with 5 layers of interconnects which include the power and ground distribution. Heat is dissipated from the back side of the MCM.

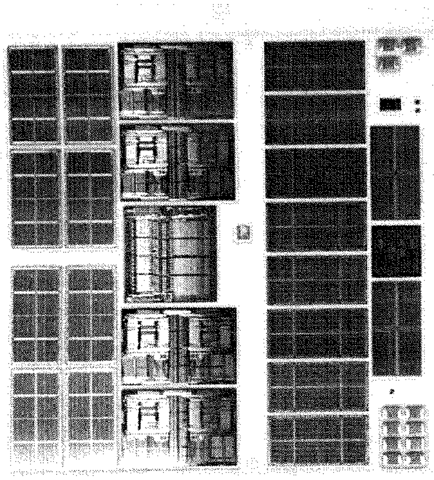

Figure 3: MCM3A

A new version of the MCM3, "MCM3B" is now in fabrication. It is similar to MCM3A except that it adds two additional layers of GE HDI. The top layer (layer 7) is a Land Grid Array (LGA), shown in figure 4a, that will be used for the I/O interconnects in place of the perimeter bond pads used in the MCM3A. The 6^{th} layer is used to interconnect layer 5 (the old top layer) to the new top layer. This new interface was added to support HIPP packaging, which will be described later.

The MCM3B's LGA is contacted using an interposer similar to the one shown in figure 4b. The interposer has conductive compressible contacts that make a connection between two mirror image LGA surfaces. This allows easy testing at the MCM3 level on a standard PCB, and later insertion into the higher level packaging scheme.

After testing, up to four MCM3Bs can be stacked one on top of the other. Their contacts to the LGA also run to two sides of the MCM package (not shown). The signals are connected to feed throughs in the sides of the ceramic package. Connection is then made up the side of the MCM using a single layer of GE HDI. The top MCM (with its exposed LGA), has the routing needed to take the signals from the lower 3 (buried) MCM3Bs to the appropriate LGA pad. Other than being thicker than single layer MCM3Bs, stacks have the same electrical interface, an LGA.

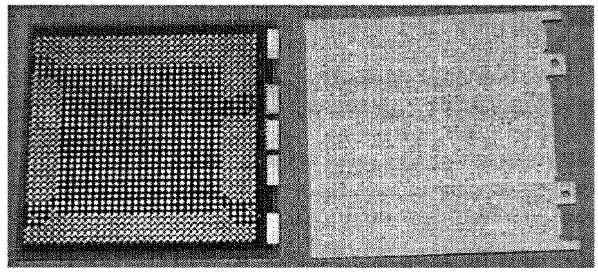

Figure 4a **Figure 4b**

Highly Integrated Packaging and Processing (HIPP):

HIPP is an AFRL program to develop efficient packaging of MCMs and various other components including standard PCBs. [4] Each MCM or PCB with its mounting hardware and LGA is referred to as a segment.

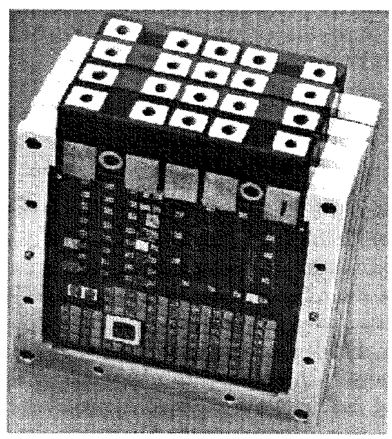

Figure 5: Highly Integrated Packaging and Processing (HIPP)

This packaging scheme brings the power in through the terminals shown at the top of fig 5. The signals are routed down a flexible printed circuit referred to as a flap. The connection between the flap and the MCM3B is made by placing an interposer between the two and compressing the contacts. One end of the flap has a ball grid array that allows the signals to make contact with a micro-back-plane. The micro-back-plane is another flexible printed circuit that interconnects the signals from various flaps. Non MCM components can be attached to a PCB which has a LGA formed on one side and the components on the other. These two by two inch PCBs can then be inserted into the HIPP structure.

Fusion Processor Segments:

In addition to 12 MCM3Bs packaged as three four-layer stacks the FP contains a 16-port Myrinet crossbar switch segment, a clock segment, and a management segment.

The Crossbar segment consists of a Myricom 16-port crossbar packaged part mounted on a two by two inch PCB with a clock and other support components. It provides non-blocking crossbar switching for internal and external communications using the Myricom System Area Network hardware protocol. This allows communication at 160 Mbytes per second in each direction for each of the 16 channels, for distances up to 10 feet.

The clock segment provides the IOBus and Myrinet clocks for the MCM3Bs. It consists of packaged parts on a two by two inch PCB.

The Management segment boots the FP. First it provides some low level boot functions by generating boundary scan signals. Then the programs are loaded by transferring data from it's Flash RAM to it's FI32 then over the Myrinet to the appropriate MCM3B. After loading each processor's memory the Management segment turns the processors on. The Management Segment's hardware is described below.

Other Segments:

AFRL's Space Vehicles Directorate has developed several other MCM based segments for Discrimination Interceptor Technology Program (DITP). These include one called the Malleable Signal Processor (MSP), and one called the MSP Management Segment. The two are usually used as a pair. The MSP contains 2 Altera 10K100A Field Programmable Gate Arrays (FPGAs) and memory. The MSP performs much of the integer processing before it's associated Management segment passes the information over the Myrinet to the FP. The Management segment contains a FI32, an Altera 10K100A, 4 Mbytes of Flash RAM, and an 8051 microprocessor with 128K of RAM and 128K of ROM. [5] [6] [7]

Sensor and Fusion Engine (SAFE):

Figure 6 shows an artist concept of the SAFE, ready for insertion into its thermal housing. It includes the fusion processor and other digital and analog circuits required for DITP. The FP alone is expected to be slightly over two by two inches and two inches long without thermal management. In Figure 6 the thermal management is provided by a phase change material. The maximum power consumption for the FP is approximately 100 Watts. For bench operation the heat can be removed with a heat sink on one side. The system is easily customized by changing the micro-back-plane to add or remove segments. In addition to the MCM segments that are available custom two by two inch PCBs are easily added.

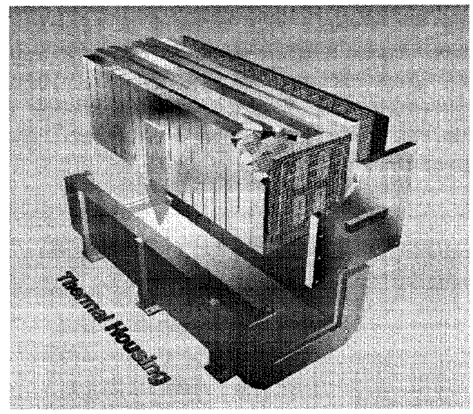

Figure 6: Sensor and Fusion Engine (Artist Conception)

Software Environment:

The GNU toolset provides the compilers (C, C++, Fortran), assembler and debugger for the FP. [8] Its operating system is the US Army developed Real-Time Executive for Microprocessor Systems (RTEMS). RTEMS provides a real-time multi-tasking operating system with open source code. [9] A simulator/communicator is available for PCs running Windows NT or LINUX, and Sun Workstations running Solaris. It simulator provides visibility into the processor operations and assists in communications with the processors for code debugging.

References:

[1] Dr. R. Linderman, Dr. M. Linderman, R. Kohler, Maj. J. Comtois, PhD, J. Sabatini "HDI Design of a 100 MFLOPS/Watt Floating Point DSP" GOMAC 1998 pg 335-338

[2] Capt. Figie, Y. Kinashi, B Johnson, R. Linderman, J. Fabunmi "Discrimination Interceptor Technology Program (DITP): Sensor Fusion for Improved Interceptor Seekers" 1996 AIAA/BMDO Missile Sciences Conference

[3] J. Rooks "Myrinet ARFL Packets Protocol", Version 1.2, 6/11/99, available from rooksj@rl.af.mil

[4] J. Lyke and G. Forman "Microengineering Aerospace Systems" H. Helvajian editor, The Aerospace Press 1999, Chapter 8

[5] J.C. Lyke, G.W. Donohoe, N. Anderson, "Malleable Signal Processor: A General-purpose Module for Sensor Integration", to appear, MAPLD 2000.

[6] G.W. Donohoe and J.C. Lyke, "Adaptive Computing for Space", Proc. Midwest Symposium on Circuits and Systems, Las Cruces, NM, August 8-11, 1999.

[7] G.W. Donohoe and J.C. Lyke, "Embedded, Configurable Computing for Space", Invited Graduate Seminar, Texas A&M University Analog and Mixed Signal Center, October 8, 1999.

[8] GNU web site www.gnu.org

[9] RTEMS web site www.rtems.com