ICOT Technical Memorandom: TM-1159

TM-1159

A VLSI Chip Set for a Large Scale Parallel Inference Machine : PIM/m

by
H. Machida, H. Ando, K. Yaduda, K. Furutani,
H. Nakashima, Y. Takeda, K. Nakajima
& M. Nakaya (Mitsubishi)

February, 1992

ICOT
Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome
Minato-ku Tokyo 108 Japan
(03)3456-3191～5
Telex ICOT J32964

Institute for New Generation Computer Technology

A VLSI Chip Set for a Large Scale Parallel Inference Machine: PIM/m

Hirohisa Machida, Hideki Ando, Kenichi Yasuda, Kiyohiro Furutani, Hiroshi Nakashima, Yasutaka Takeda*, Katsuto Nakajima*, and Masao Nakaya*

LSI Laboratory, Mitsubishi Electric Corporation
4-1 Mizuhara, Itami, Hyogo, 664 Japan
Phone : 0727-84-7347 Fax: 0727-84-7439

**Computer & Information Systems Laboratory, Ofuna, Japan*

Abstract - Three VLSI chips which are a processor chip, a cache memory chip, and a network control chip for a highly parallel inference machine with capability of max 128 MRPS (Mega Reduction Per Second) have been developed. A processing element (PE) which consists of the processor chip and the cache memory chip has been designed to be suited for logic programming languages. The processor chip has been constructed in a submicron CMOS process technologies. The cache memory chip implements a hardware support called "Trail Buffer" which is suitable for the execution of the Prolog-like languages. The network control chip makes it possible to connect 256 PEs in a mesh network.

I. Introduction

A high performance machine has been required for a massive knowledge information processing. It is not suitable for general purpose computers to make an inference from huge amounts of data. In order to solve this problem, expert systems which execute logic programming languages in high speed have been developed. To attain higher performance, it is effective to make a lot of processing elements operate in parallel and to reduce the machine cycle time of processing element.

A large scale Parallel Inference Machine (PIM/m) is being developed in the Japanese Fifth Generation Computer Systems Project [1]. The maximum configuration of PIM/m is shown in Figure 1. Up to 256 processing elements are connected to form a mesh network. The processing element has the processing unit, the cache unit, the main memory, the floating point processor, and the network control unit as shown in Figure 1. We have developed a processor (PU) chip, a cache memory (CU) chip, and a network control (NU) chip required for high speed operation and for highly parallel processing. All chips are packaged in 361 pin-PGA and operate by a single 5-V power supply. Total power consumption of three chips is 7 W at 16.7-MHz operation. All device features are summarized in Table 1.

The PU chip and the CU chip are also key components of an AI workstation. The PU chip has capability to execute two different type logic programming languages, KL1[2] for PIM/m and ESP[3] for the AI workstation. KL1 is a parallel logic programming language and is very powerful to represent parallel process communicating. ESP is the language in which Prolog and object oriented language features are combined. The NU chip has four bidirectional channels to connect adjacent four PEs and two buffers for message packets. The packet transmission and buffering are automatically performed without any interruption of the execution of the PU and the CU chip.

This paper describes characteristics of these three chips that can be used to develop a large scale parallel inference machine : PIM/m.

II. Processor Chip

The PU chip is a 5-stage pipelined microprocessor under the control of a microprogram stored in a 32K-word control store [4] and this chip was introduced in CICC'91 [5]. The PU chip supports fast execution of the logic programming languages through a tagged 40-bit architecture. A 40-bit data includes an 8-bit tag which indicates a data type.

The block diagram of the PU chip is shown in Figure 2. The performance of the PU chip has been improved by a factor of 3.78 as compared with the previous non-pipelined processor. Achieving the high performance, this chip has powerful mechanism of pipelining, tag manipulation, and 2-stage clock scheme of layout and is fabricated in a 0.8μm double-metal CMOS technology.

The execution mechanisms of the logic programming languages stand on unification, data typing and dereference are very important operations for efficient implementation. For these operation, the PU chip has powerful mechanisms to manipulate tagged data. Especially, the pipelined data typing and dereference are the most unique features. Those mechanisms greatly contribute to the high performance with a capability of 0.5 MRPS/PE in append operation.

III. Cache Memory Chip

Figure 3 shows the block diagram of the CU chip, which contains of an instruction cache block, a data cache block, and a main memory interface block as a DRAM controller [6]. This chip includes the "Trail Buffer" with 16-word x 32-bit RAM which accelerates the execution of the logic programming languages. The detailed cache features are listed in Table 2.

A. Cache configuration

The instruction and data cache blocks employ the physical address caching scheme with on-chip translation lookaside buffers (TLB's) for high hit-ratio. We have determined the cache memory size by practical simulation taking the relationship between the chip size and hit-ratio of the cache memory into configuration. Figure 4 and Figure 5 show the simulation result of hit-ratio versus the TLB size and cache size, respectively. Hit ratios of 99.83% and 99.88 % are obtained by the 32-entry, 2-way set associative TLB's for the instruction and data address translations. And hit-ratios of 94.4% and 99.2 % are obtained with the cache memories, respectively. The scan test method and the special command to access every memory cell are applied to enhance the testability.

B. Support for logical inference

To accelerate logical inference, the CU chip has hardware called "Trail Buffer". In Prolog-like languages, there is an operation called "Backtrack". When backtrack operation occurs, a variable assignment by unification should be reset. In this operation, when a processor unifies a variable, the address of the variable should be stored in "Trail Stack". The trail buffer, which can hold up to 16 addresses of assigned variables, is a cache of the trail stack.

The functions of the trail buffer are to store the addresses of assigned variables, to supply address data of the variable to be reset, and to remember the number of reset times. It is estimated by simulation, that the trail buffer improves the speed of the logical inference by 10% (average).

IV. Network Control Chip

In a loosely coupled highly parallel computer system, it is important to develop network circuits which realize high speed communication and reduce loads of processors.

A. Hardware

Figure 6 shows the block diagram of the NU chip. It consists of a network control unit, a trace memory, a system timer, a maintenance control unit, and an external bus control unit. The network control unit has five pairs of channels which are used for four adjacent PEs and its own processor chip. The channels which are connected to interface bus have 1024 x 9-bit read buffer memory (RB) and 1024 x 9-bit write buffer memory (WB). Each of four transmission channels has 64 x 10-bit FIFO buffer memory to reduce the possibility of network choking. Packet data are transferred asynchronously in 10-bit parallel including a parity bit. The external bus control unit has a function which arbitrates among the NU chip, FPP, and SCSI.

B. Switch circuits

The NU chip supports a message-passing communication between PEs. Switch circuits are shown in Figure 7. Four adjacent PEs and its own processor unit are connected by 5 x 4 switch circuits. Each receiving channel has its own path table (PT) to determine the channel to transmit a packet. Channel controller looks up this table using the destination address of the packet, and connects receiving and transmission channels according to the PT data. Packets are, therefore, switched in parallel regardless of the processor operation. The index of the PT is the coordinate of the destination point which is represented by a 1024 x 2 bit map. Each entry of the PT contains the channel number to which packets are transmitted. The network control scheme is summarized in Table 3.

V. Performance Evaluation

Figure 8 shows a typical schmoo plot of cycle time versus supply voltage of the NU chip. This chip achieves 80-Mbits/sec data transmission for each channel at 33-MHz operation.

These three chips have been successfully constructed. PIM/m which is utilized by these chips has passed system level tests, and has correctly executed the KL1 at 15.4 MHz under worst case condition. It has peak performance of 128 MRPS (Mega Reduction Per Second) and 40-Mbits/sec data transmission for each channel. Table 4 shows a performance comparison with a prototype parallel inference machine [7] which used eleven gate arrays, memories, and glue logics. Owing to these three VLSI, PIM/m achieves about 15 times higher than the prototype machine at peak performance. Though the number of PE in PIM/m is 4 times as many as previous one, it keeps the same size. The power consumption increase of PIM/m is only 30 %.

VI. Conclusion

A VLSI chip set for a high performance large scale parallel inference machine has been developed. The die photos of these three chips are shown in Figure 9. All chips operate by a single 5-V power supply. Total power consumption of three chips is 7 W under a typical condition. A photograph of a 32.5-cm x 31.0-cm PIM/m processing element board is shown in Figure 10. These chips make it possible to construct a high performance, small size, and low power consumption parallel inference machine.

ACKNOWLEDGMENT

The authors wish to thank Dr. H. Komiya, Dr. T. Nakano, and Dr. Y. Horiba for their encouragement and support of this research program, and also thank Dr. S. Uchida and researchers in ICOT for giving us the opportunity to conduct this research.

We also greatly indebted to the device processing staff members who have contributed to the development of these devices, and to Mr. H. Kakiuchi for the testing.

REFERENCES

[1] Uchida, S., Taki, K., Nakajima, K., Goto, A., and Chikayama, T.,"Research and Development of the Parallel Inference System in the Intermediate Stage of the FGCS Project," in Proc. of Intl. Conf. on Fifth Generation Computer Systems 1988, pp. 16-36, 1988.

[2] Chikayama, T., Sato, H., and Miyazaki, T. "Overview of the Parallel Inference Machine Operating System(PIMOS)," in Proc. of Intl. Conf. on Fifth Generation Computer Systems 1988, pp.16-36, 1988.

[3] Chikayama, T. "Unique Features of ESP," in Proc. of Intl. Conf. on Fifth Generation Computer Systems 1984, pp. 292- 298, 1984.

[4] Nakashima, H., Takeda, Y., Nakajima, K., Andou, H., and Furutani, K., "A Pipelined Microprocessor for Logic Programming Languages," in IEEE Proc. of Intl. Conf. on Computer Design, pp. 355-359, 1990.

[5] Machida, H., Ando,H., Ikenaga,C., Maeda,A., Nakashima, H., and Nakaya, M., "A 1.5 MLIPS 40-bit AI processor," in Proc. of the IEEE 1991 Custom Integrated Circuits Conf., pp.15.3.1- 15.3.4, 1991.

[6] Yasuda, K., Furutani, K., Maeda, A., Wakano, S., Nakashima, H., Takeda, Y., and Yamada, M.," An Intelligent Cache Memory Chip suitable for Logical Inference," IEICE Tran. on Electronics, Vol. E 74, No. 11, pp. 3796-3802, 1991.

[7] Takeda, Y., Nakashima, H., Masuda, K., Chikayama, T., and Taki, K., "A Load Balancing Mechanism for Large Scale Multiprocessor Systems and Its Implementation," in Proc. of Intl. Conf. on Fifth Generation Computer Systems 1988, pp. 978-986, 1988.

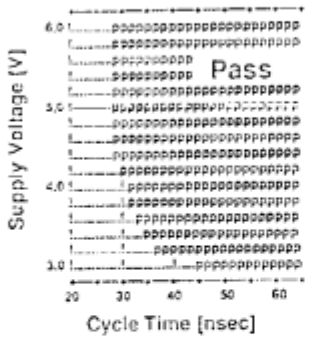

Figure 8 Schmoo Plot of Network Control Chip

Table 1 Device Features

	Processor chip	Cache memory chip	Network control chip
Chip size	16.3 x 13.6 mm	14.5x14.8mm	14.2x14.0mm
Transistor count	984K	610K	329K
RAMs	274K	545K	251K
Logics	110K	65K	68K
Scan path stage	409	420	268
Pad count	352	347	325
Package	361 pin PGA	361 pin PGA	361 pin PGA
Chip cycle time	30nsec	33nsec	28nsec
Power supply	5V	5V	5V
Power consumption	2.5W	2.3W	2.2W
Process technology	0.8μm double-metal CMOS	1.0μm double-metal CMOS	1.0μm double-metal CMOS
System clock	16.7MHz (typical condition)		
System performance	128 MRPS		

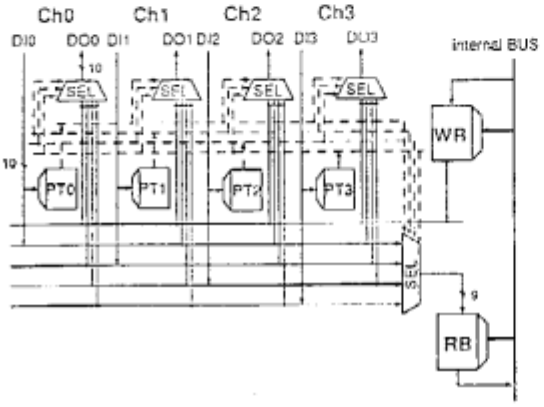

Figure 7 Block Diagram of Switch Circuits

Table 2 Cache Features

	Instruction cache	Data cache
Address	Physical	Physical
Associativity	Direct map	Direct map
Cache size	5KBytes	20KBytes (off chip)
Tag entry	256	1024
TLB	32entry 2-way	32entry 2-way
TLB replacement	LRU	LRU
Parity		SECDED
Consistency		write back DRAM refresh

Table 3 Characteristics of NU chip

Inter-PE Communications	
Operation Mode	Asynchronous
Control Strategy	Distributed Control
Switching Methodology	Packet Switching
Network Topology	Mesh Connection
Number of Channels	5 (4 : adjacent, 1 : own)
Width of Channel	10 bits (9 bits : data, 1 bit : parity)
Transmission Buffer	640 bits
Read Buffer	9 Kbits
Write Buffer	9 Kbits
Bandwidth	40 Mbits/sec

Table 4 Performance Comparison

	Number of PE	cycle time [nsec]	Peak performance [MRPS]	Capacity [mm³]	Power consumption [KVA]
prototype	64	200	6.64	9 x 960 x 780 x 1400	24.0
PiM/m	256	60	128	8 x 960 x 780 x 1400	31.2

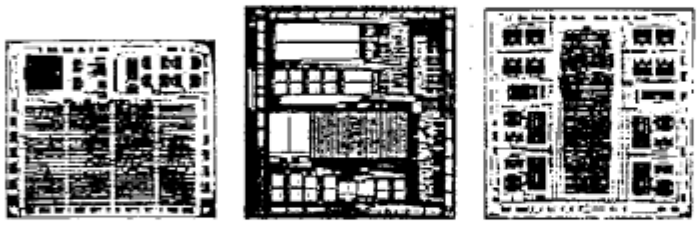

(a) Processor chip (b) Cache memory chip (c) Network control chip

Figure 9 Chip Photomicrograph

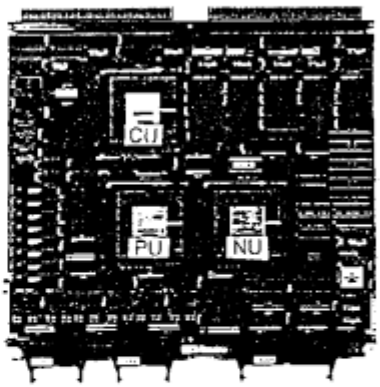

Figure 10 Processing Element Board

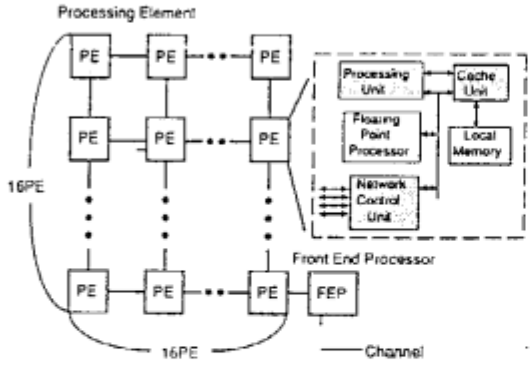

Figure 1 Configuration of PIM/m

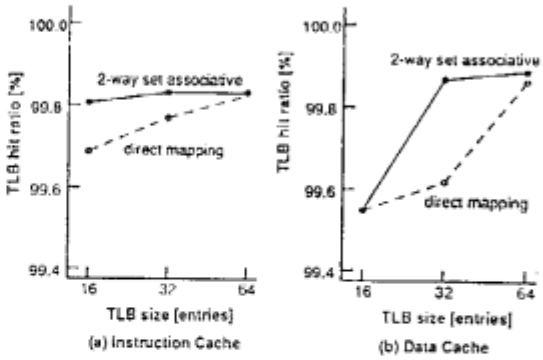

Figure 4 Simulation Result of TLB Hit Ratio

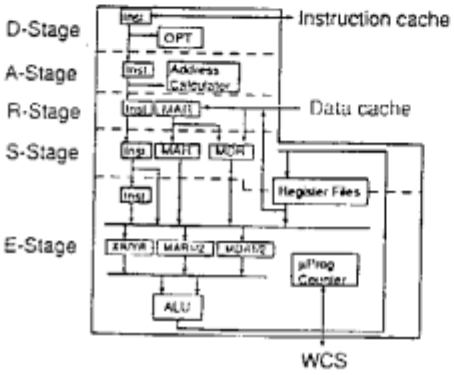

Figure 2 Block Diagram of Processor Chip

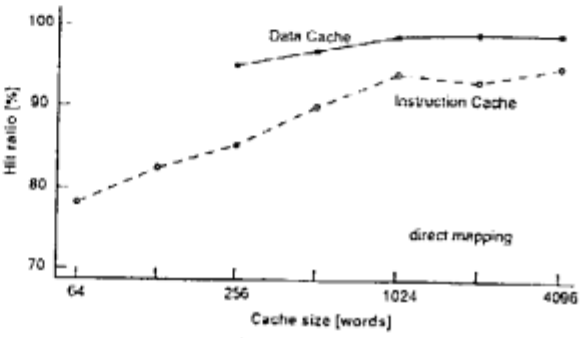

Figure 5 Simulation Result of Cach Hit Ratio

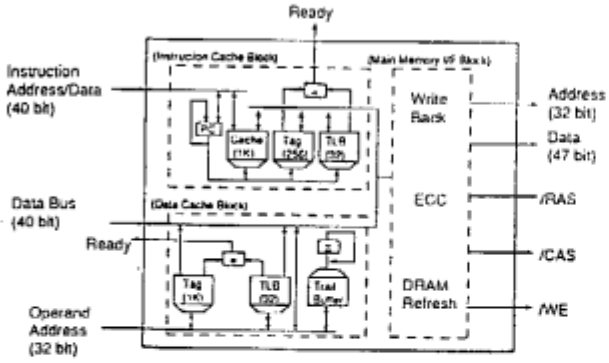

Figure 3 Block Diagram of Cache Memory Chip

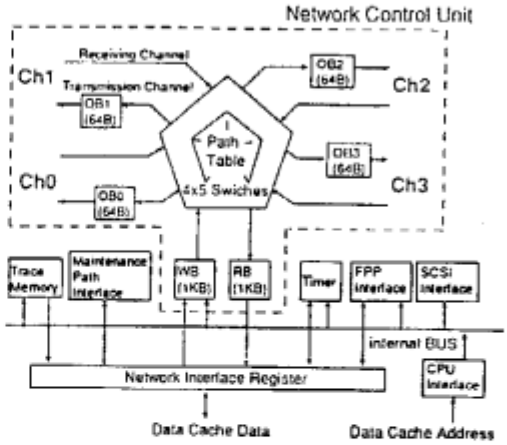

Figure 6 Block Diagram of Network Control Chip