LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

LA-UR-84-1324
CONF

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

CONF-8405131--5

TITLE: CURRENT RESEARCH IN PARALLEL MICROPROCESSING SYSTEMS AT LOS ALAMOS

LA-UR--84-1324

DE84 011342

AUTHOR(S): C. DWAYNE ETHRIDGE

SUBMITTED TO: IDEAS IN SCIENCE AND ELECTRONICS

DISCLAIMER

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

CURRENT RESEARCH IN PARALLEL MICROPROCESSING SYSTEMS AT LOS ALAMOS

C. Dwayne Ethridge
Los Alamos National Laboratory
Computing and Communications Division
P.O. Box 1663 MS-B255
Los Alamos, New Mexico 87545
(505) 667-3310

Abstract

The Computing and Communications Division at the Los Alamos National Laboratory has designed and is building a parallel microprocessor system (PuPS) to serve as a research tool for evaluating parallel processing of large-scale scientific codes. PuPS is an experimental architecture consisting of an orthogonal array of 20 processing elements by 32 memory elements, establishing a tightly coupled, shared-memory (16-Mbyte) machine. The hardware incorporates VLSI components, such as 16-bit microprocessors, floating-point co-processors, and dynamic random access memories. The design replaces conventional MSI/SSI circuitry with programmable array logic, logic sequencers, and logic arrays. This experimental system, which is only 1 element of the parallel processing research being done by the Laboratory's Computing and Communications Division, will enable direct comparisons of speedups of algorithms for a variety of multiprocessor architectures.

Introduction

The objective of the parallel processing research at the Los Alamos National Laboratory is to assess the potential of parallel processing to provide increases in speed for large-scale computations in areas such as hydrodynamics, Monte Carlo techniques, reactor safety studies, and nuclear waste disposal. Speedup in parallel processing (the ratio of a computation execution time using 1 processor to the computation time using p number of processors) is related to the code parallelism and to the number of processors. There are no processor-rich systems (16 or more processors) at this time, although computer manufactures may be able to offer such systems by 1990; furthermore, there are no processor-rich systems in the U.S. with floating-point arithmetic, large memory, and Fortran programmability. Therefore, as part of the parallel processing research at Los Alamos, we are constructing an experimental parallel microprocessing system to serve as a research tool for evaluating parallel processing of large codes on various multiprocessing architectures.

PuPS is an experimental computer design consisting of an orthogonal array of 20 processing elements by 32 memory elements establishing a tightly coupled, shared-memory (16-Mbyte) machine. The system will have reconfigurable processor-to-memory and processor-to-processor interconnections to enable direct comparisons of speedups of algorithms for a variety of parallel architectures. The reconfigurable design allows architectures with multiple processors sharing common memory and with networks of processors with only local memory, such as rings, trees, and stars.

The hardware design incorporates VLSI components, such as 16-bit microprocessors, floating-point co-processors, dynamic random access memories (DRAM),

and replaces conventional MSI/SSI circuitry with programmable array logic (PAL), programmable logic sequencers (PLS), and programmable logic arrays (PLA). Use of these arrays greatly reduces the different types of logic devices in the system, which in turn reduces engineering effort related to parts procurement as well as in maintaining parts for computer support and maintenance after the computer is in operation.

The design goals for the experimental parallel microprocessor include the following requirements: processor-rich (more than 8 processors), large memory, commercially available components, floating-point hardware, Fortran compilation, and nominal software development. This experimental parallel microprocessor system is 1 element of the parallel processing research within the Computing and Communications Division at the Los Alamos National Laboratory.

PuPS Architecture

PuPS is a tightly coupled, shared-memory, multiple-instruction multiple-data (MIMD) machine that supports reconfiguration between processor and memory nodes to permit experimentation with common memory architectures and with various processor network structures, such as rings, trees, and stars. The proper combinations of processors and memories can be selected based on the parallel processing objectives. The system consists of numerous processor and memory nodes that are directly interconnected using multiple processor-to-memory buses and multiported global memory nodes. The multiple-bus/multiported memory design functionally implements a full crossbar switch between the processor and memory nodes.

This multiple-bus architecture allows processor-to- processor communications to occur concurrently with processor execution from either local memory or global memory. Reconfiguration between processor and global memory nodes is accomplished through a memory mapping facility within each processor node. Reconfiguration will be limited to initialization time only; however, experience with the hardware and software could lead to dynamic reconfiguration if applicable.

The system architecture includes the PE, which is the general processing element or execution processor, and the GM, which is the global memory element. Each processing element (PE) has local memory (LM) that consists of both programmable-read-only memory (PROM) and random access memory (RAM). A control processor (PC) loads the execution code into the global memories and initiates execution through the system control bus. Systems with a large common memory can be configured with the processing elements and global memories.

Systems with only local memory can be configured with the general processing elements, global memories, and data transfer processors. A data transfer processor consists of a read processor (RP), which uses 1 bus to read from memory and to store in a first-in/first-out (FIFO) memory, and a write processor (WP), which uses 1 bus to read the FIFO and to write to memory.

This architecture is enabled by the orthogonal packaging scheme. The processing element cards plug horizontally into the processor section of the backplane. The global memory element cards plug vertically into the left and right global memory sections (low- and high-memory addresses) of the backplane. The left and right sections provide contiguous logical addressing. However, the left and right sections are electrically isolated, limiting each backplane section to 16 global memory cards and 1 bus termination card. Therefore, the backplane sections provide 21 microcomputer buses: 1 control processor bus and 20 processing element buses. The control processor bus is Multibus.

Each processing element bus (backplane bus) consists of 32 lines. Twenty-four lines provide 24 address signals, 16 of which are multiplexed address/data signals. Seven lines provide bus control and memory error identification. The remaining line is a reserved line, possibly for a diagnostic bus.

Processing Element Implementation

Selection of the floating-point hardware for the processing element was required early in the design cycle. The design goal of using commercially available components led to the selection of the only available floating-point co-processors at that time: the Intel iAPX 86/20, which consists of the Intel 8087 and 8086 co-processors. The Intel 8087 is the high-performance, numeric data co-processor for the processing element or execution processor. This co-processor provides the floating-point hardware required to create a processing element with significant scientific computational capabilities. The Intel 8086 16-bit HMOS Microprocessor allows the processing element to directly address a Mbyte of global memory. Memory mapping elements 74LS610s used in the design extend addressing to 16 Mbytes.

Three types of processor nodes are included within the system: the system control processor, the processing elements, and the data transfer processors. The system control processor performs system initialization (downloading of global memory, configuration control, etc), initiation of parallel processing applications code, performance measurements, and memory error processing. This processor also incorporates the interprocessor interruption facility. In addition, because the multiprocessor is strictly an execution environment, the system control processor provides communication with an external local area network that includes development workstations. Each processing element includes the Intel 8086/8087 co-processors, 8 to 32 Kbytes of local dedicated PROM and 4 to 16 Kbytes of RAM, real-time interrupt facility, and memory mapping logic that allows sixty-one 16-Kbyte segments to be permanently and/or dynamically allocated within the system global memory. Each data transfer processor is a high-speed controller specifically designed for implementing processor-to-processor communications by performing data movement between global memory segments. Two Intel 8089 HMOS Input/Output Processors are used to implement each data transfer processor.

Global Memory Element Implementation

The system global memory consists of multiple memory nodes, each having a 256- to 512-Kbyte RAM array (Microbar DBR50-256) accessible from the system control processor, and a multiported memory controller. The port for the system control processor supports downloading and memory error reporting functions. The multiported memory controller includes interface logic for 20 ports, memory arbitration logic that implements a last-granted/ lowest priority algorithm, and a high-speed memory access controller. Memory mapping logic within each processor node allows each memory node to be allocated as either private or public memory for each processor node.

All 32 global memory cards are identical in construction. The upper 5 address bits (minus the most significant bit used by the left/right bus controllers) are compared with the 4-bit backplane geographical address to grant memory access on the specific card.

The processor-memory interconnection is accomplished with memory mapping logic at each processor node, a multiported memory controller at each global memory node, and a multiple bus interconnection backplane that allows an orthogonal arrangement of processor and global memory boards. The packaging scheme uses minimal bus lengths in

providing complete physical interconnection between the processor and global memory nodes. The processor-memory interconnection provides a fully reconfigurable processor-to-memory connection, resolves access arbitration when multiple processors are simultaneously accessing a common global memory node, and supports mutual exclusion to shared memory. Control of shared memory is accomplished through an extension of the lock mechanism available with the Intel 8086/8087/8089 co-processors.

Backplane Bus Control

A memory cycle (global memory) using the system backplane is controlled by 5 hardware controllers. An Intel 8288 Bus Controller provides the microprocessor bus control. A left or a right backplane bus controller located on the execution processor provides the bus control signals. A port interface controller (1 for each bus, 20 per global memory card) provides the control signals (P0 through P19) and initiates a request to the last-granted/lowest priority bus arbitration controller. This controller issues a grant to only 1 of the 20 port interface controllers. The final controller, the high-speed bus controller, multiplexes the address and data bus connection to the Microbar Memory board. The principal hardware components of these controllers are Monolithic Memories programmable array logic devices (PAL16L8) and programmable logic sequencers (PAL16R8).

PuPS Hardware and Software Development

The following equipment is involved in the development of PuPS. The microcomputer development system MDS is an Intel Series 3 MDS-800 with in-circuit-emulation capability ICE-86A for the iAPX 86/20 co-processors. The MDS system includes a Data I/O System 19 Universal Programmer for the local memory PROMs and for the programmable logic devices. Memory storage for the system is provided by a 40-Mbyte hard disk. System files and source file entry and/or backup are provided by a single, double-density floppy disk drive.

A local area network links the parallel microprocessor, the MDS, the backup storage (VAX 780) and the software workstation. The software workstation will be used to convert existing Fortran production code to parallelized, compiled, linked, and located code for the execution processor. The operating system for the software workstation and possibly the control processor will be Intel's iRMX 86 Operating System, which provides a multiterminal and multiuser interface capability.

General System Construction

The PuPS enclosure is approximately 1 meter deep, 2 meters high, and 3 meters wide. The MDS includes a 40-Mbyte hard disk unit. The processor card cage section provides the housing for 27 processor cards with individual circuit breakers. Seven card slots are reserved for the control processor leaving 20 card slots for the processing elements.

Two 200-ampere power supplies provide 5 volts to 1 global memory section. One 200-ampere power supply provides 5 volts to the processor section. Voltage (and return) is applied to each card at the front edge so that no supply voltage is in the backplane wiring. Each card voltage is also switched by its own circuit breaker.

Summarizing Remarks

The Parallel Microprocessor System is an experimental computer architecture design consisting of an orthogonal array of 20 processing elements by 32 memory elements establishing a tightly coupled shared-memory (16-Mbyte) machine. The principal objective is to develop an experimental parallel microprocessor system to serve as a research tool for evaluating parallel processing of production codes on various multiprocessor architectures.

The hardware design incorporates VLSI components, such as 16-bit microprocessors, floating-point co-processors, dynamic random access memories, and replaces conventional MSI/SSI circuitry with programmable array logic, programmable logic sequencers, and programmable logic arrays. This experimental parallel microprocessor system is 1 element of the parallel processing research within the Computing and Communications Division at the Los Alamos National Laboratory.

Acknowledgments

The work is being performed under the auspices of the U.S. Department of Energy. Reference to a company or product name anywhere in this paper does not imply approval or recommendation of the product by the University of California or the U.S. Department of Energy to the exclusion of others that may be suitable.