

UDK 681.3.02

Saša Prešern
Iskra Delta and
Jozef Stefan Institute, Ljubljana

ABSTRACT. This paper is a selected survey of parallel computer systems. A classification of parallel computers is given and some most attractive architectures are discussed. Special attention is paid to massively parallel processors. The organization and interconnection structure of multiprocessor systems is given. By analysing a trend of research in parallel computer systems over last 10 years some predictions are given about individual features which will probably have great influence on future parallel computer systems. An extensive survey of references in parallel computer systems is given.

IZBOR IN PREGLED PARALELNH RACUNALNIŠKIH SISTEMOV. Članek podaja izbor in pregled paralelnih računalniških sistemov. Narejena je klasifikacija paralelnih računalnikov in opis nekaterih najbolj zanimivih arhitektur. Podana je organizacija multiprocesorjev in opisane so različne povezovalne strukture med procesorji ter pomnilniki v posameznih sistemih. Analiza trenda raziskav paralelnih računalniških sistemov v zadnjem desetletju omogoča izločitev posameznih značilnosti, ki bodo predvidoma močno vplivale na razvoj bodočih paralelnih računalniških sistemov. V bibliografiji je priložen obširen pregled referenc za paralelne računalniške sisteme.

1. INTRODUCTION - EVERYBODY MAKES IT PARALLEL

A few years ago all high developed countries in the world have started projects in developing a parallel computer system. All these projects were financially supported by governments. Many companies and research institutes also started research projects on parallel systems. The falling price of microcomputers and VLSI facilities on universities has encouraged many universities to design and to build parallel computer architectures based on linking many microprocessors or specially designed VLSI chips together to work on one job.

Development of a parallel computer is an extremely difficult task which includes:

- development a new concept of parallel computer architecture,
- design of an operating system that supports parallel architecture,
- transformation of traditional sequential application programs to parallel programs either by preprocessor or by a parallel programming language.

We see that by switching from SISD (single instruction single data) machines to MIMD (multiple instruction multiple data) machines one can not simply upgrade an existing SISD computer system but one is faced with problems which are conceptually new. Research and development of a parallel computer system requires a very strong research which often includes:

- more than 100 specialists,
- a billion dollar financial support,
- research and development phase which lasts several years.

Government financial support is only a fraction of the whole finances which are devoted to projects in parallel computing. Strategy makers in most companies are familiar with market research studies which predict that parallel processing machines will take about 50 percent of the market in high-performance computers by 1990.

2. CLASSIFICATION OF PARALLEL SYSTEMS

Parallel computers are usually divided into three architectural configurations:

- SIMD pipelined computers
 - * early vector processors,
 - * attached processors,
 - * recent vector processors,
 - * other vector processors,
- SIMD array processors,
- MIMD parallel processors,
 - * massively parallel processors,
 - * small scale parallel systems.

Another grouping is possible as for example classification according to distribution of local and global memory into tightly and loosely coupled parallel systems or classification according to application possibilities into general purpose or special

purpose computers.

Many existing computers are now using several parallel approaches. Parallelism in pipeline computers is performed by overlapping computations and is therefore temporal parallelism. Parallelism in array processors is performed by multiple synchronized ALUs and is therefore spatial parallelism. Parallelism in multiprocessor systems is performed by a set of processors with shared resources which work in asynchronous mode.

The list of projects in parallel computing is getting longer every day. By comparing the architectural approach in different projects we see that the computer scene in parallel computer systems is particularly varied. It is difficult to classify parallel computers, but helpful in order to concentrate on similarities and differences between the computer architectures. Because parallel computers are using several different architectural principles one might argue a proposed classification.

Some of described computer are "paper machines" that have been studied theoretically and by simulation, but have not been build. Many of this projects were funded by government agencies, but some of them are industry projects (IBM, Burroughs, CDC,...).

There follows an alphabetic list of the parallel computer systems or projects, each with the name of the chief architect and host institution. A list of references dealing with each project is also given. The most interesting architectures are briefly described. The list of parallel computers is grouped according to upper classification.

SIMD PIPELINE COMPUTERS

=====

EARLY VECTOR PROCESSORS

BVM (Boolean Vector Machine), Robert A. Wagner, Duke University, North Carolina. This is a collection of 1-bit processing elements connected as a hypercube with rings at each corner, using the Cube-Connected-Cycle topology.

STAR-100, Control Data Corporation. The design of Star started in 1965 and was delivered in 1973. This is a processor with two nonhomogeneous arithmetic pipelines. (HWAB5, LINB2, PUR74).

TI ABC, (Texas Instruments Advanced Scientific Computer), Texas Instrument. This machine uses 1 to 4 homogeneous pipelines and was delivered in 1972. (HWAB5, KOB81).

ATTACHED PIPELINE PROCESSORS

CSPI MAXIM/64, (CSP Inc., BILLERICA, Massachusetts). Maxim/64 in a minimal configuration includes a 16 slot chassis, a 64-bit floating point array processor, 16 Mbytes of data memory and Micro VAX-II CPU. The machine is designed for research, scientific and engineering users and costs about \$170.000. (NAN86).

FPS-AP120, Floating point systems, Beaverton, Oregon, USA. This company produces also a new version of attached pipeline processors FPS-164 and FPS-

264 which is used in configuration named LCAP (Loosely coupled array of processors). More than 1500 machines had been sold and were used mostly for signal processing. They are quite cost effective in comparison to Cray or Cyber computers. (HOC81, HWAB5, WIL82).

IBM 3838

IBM 3838 is a multiple pipeline scientific processor specially designed to attach to IBM mainframes, like the System/370, for enhancing the vector-processing capability of the host machine. It is microprogrammed pipeline processor which can be supplied with custom-ordered instruction sets for specific vector applications.

RECENT VECTOR PROCESSORS

Cray-1 Cray Research Inc., Chippewa Falls, Wisconsin, USA.

This is the first successful vector computer. More than 40 computers have been sold and installed, first in 1976. It comprises 12 special-purpose pipelines for the different arithmetic operations. It is very expensive. (HWAB5, JOR82, RUS78).

An upgrade of this computer is Cray-2, (HOL85).

Cyber-205

This computer is an example of pipelined architecture and is highly competitive with the CRAY-1. It is based on CDC STAR 100. It is based on one, two or four pipelined general-purpose units working always to and from main memory. It is an expensive machine, designed initially to weapons' calculations and weather simulation. (HOC81, HWAB5, VON84).

CDC/NASD Control Data Corporation Numerical Aerodynamic Simulation Facility.

This is a supercomputer to be used in 1990s for aerospace vehicle or superjet designs. The speed requirements was set to be at least 1000 Mflops and the purpose is to calculate the viscous Navier-Stokes fluid equations for three dimensional modeling of the wind tunnel experiments. (HWAB5, HOC81).

VP-200, Fujitsu.

This system has a scalar and a vector processor which can operate concurrently and it can be used as a loosely coupled back-end system. (HWAB5, LLU84, UCH85).

OTHER VECTOR PROCESSORS

Ahmdal 1200

This computer is a European version of Fujitsu's recent vector processor VP-200. Similar version of VP-100 is known in Europe as Ahmdal 1100 computer. (KOC85).

Siemens VP200

This is another European version of Fujitsu's vector processor VP-200. Fujitsu's VP-100 is as Siemens product known as Siemens VP100. (KOC85).

YH1

This is China's first supercomputer, known also as "Galaxy". The development started in 1978 at the University of Defense Science and Technology in Changsa. The machine looks like a Cray computer. (NEW85/1).

SIMD ARRAY COMPUTERS
 =====

BSP, (Burroughs Scientific Processor),
 Burroughs (HOCB1, HWA85, KUC82).
 BSP has been largely based on the experiences that Burroughs have gained as major contractors on the ILLIAC IV project. The design principles of the BSP were to provide a machine using a standard technology, which would be programmed in a high level language and sustain a continuous 20-40 Mflops/s.

ICL DAP (ICL Distributed Array Processor)
 (HOCB1).
 This is an array of one-bit processors which are often called associative processors. The design of pilot DAP was started in 1974 and consisted of a two-dimensional arrays of 1024 1-bit processors.

ILLIAC-IV (BAR68, DAV69, BQU72, HWA85)
 This computer was designed for the solution of partial differential equations and can be described as an 8x8 array of 64-bit floating point processing elements each (PE) with 2Kwords of memory. It was working with nearest-neighbor connections (fig. 1) and controlled by a single instruction stream processed in a central control unit.




Fig. 1.: The connectivity between 64 processing elements in ILLIAC IV (HWA85).

MPP, (Massively Parallel Processor)
 This processor was developed for processing satellite imagery at the NASA Goddard Space Flight Center and has 128x128=16384 microprocessors that can be used in parallel. Each processor is associated with a 1024-bit RAM. (BAT80, BAT82, HWA85)

PEPE (Parallel Element Processor Ensemble).
 This special purpose computer is Burroughs floating point processor array which was developed at Bell Laboratories and designed to control a ballistic missile defense system of radar detectors and missile launchers. This is a loosely coupled system of 288 processing elements, each containing three processors. (KAR82, YAW77, FIN77)

STARAN
 In this processor a bit serial associative memory is used. Staran consists of up to 32 associative array modules each containing 256 processing elements. The first Staran was installed for digital image processing in 1975. (YAW77, RUD72, BAT77, KAR82)

MIMD PARALLEL PROCESSORS
 =====

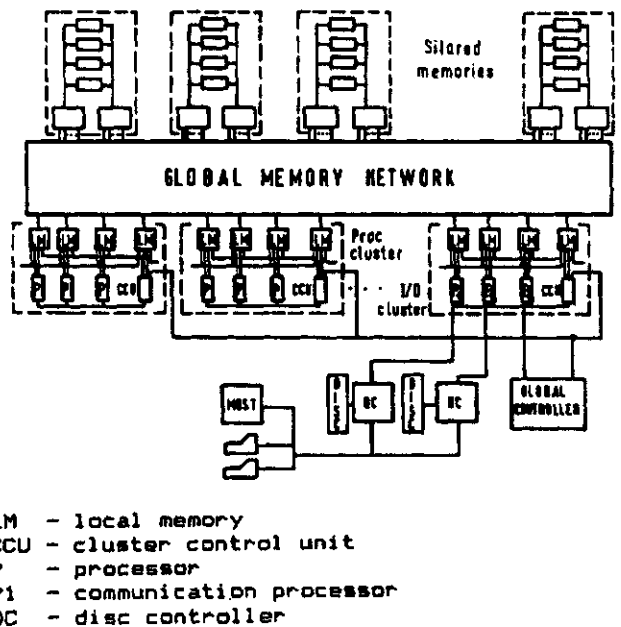
MASSIVELY PARALLEL PROCESSORS

The main accent in this architecture is in interconnection mechanism to connect several hundred processors with memory modules. A high processing power is achieved even by applying a standard processors.

BUTTERFLY, BBN - Bolt, Beranek & Newman
 (HOL85/2, HOL86, RET86)
 The Butterfly computer is a large scale shared memory parallel processor that achieves high performance in configurations as large as 256 processors. The processors used are Motorola 68020 with Motorola 68881 floating point hardware. The system has a maximum performance of 256 MIPs of processing power in 1 MIP increment and up to 1 Gbyte of memory in 4 Mbyte increments.

Processor-memory interconnection is realized via a multistage self-routing switch network. All processors can access memory simultaneously and in parallel, provided that no two processors try to take the same output path from a particular node. Butterfly network for 16 processors and 16 memories, called barrel-switching network is shown in fig 17. The speedup is nearly linear and is measured in a system with 256 processors ranging from 180 to 230 times that of a single processor (fig 22).

CEDAR, David Kuck, Duncan Lawrie and Daniel Gajski, University of Illinois at Urbana-Champaign, USA. (ABU84, ABU86).
 Cedar is an eight year project that started in 1983. The architecture is hierarchical: sixteen clusters of eight processing elements are connected via an extended Omega global switching network to 256 global memory modules of 4 to 16 Mwords each. Each cluster has eight processing elements, each with 16 kwords of local memory. These processing elements are pipelined and interconnected via a local switching network (fig 2.)



- LM - local memory
- CCU - cluster control unit
- P - processor
- P1 - communication processor
- DC - disc controller

Fig. 2.: The architecture of CEDAR parallel computer (ABU85).

The prototype Cedar 32 has four clusters of eight PEs and uses 400 ns clock period. This

gives a total maximum performance of 80 Mflops/s (Comparable to the Cray-1) for the desk-top sized prototype. Cedar 128 will have 16 clusters, giving total maximum performance of 320 Mflops (1988) and Cedar 512 will have 64 clusters, giving total maximum performance of 1.2 Gflop/s (1990). An alternative engineering is planned using 40 ns clock period giving a four cluster Cedar 32H 800 Mflops/s (1989) and 16 cluster Cedar 128H 3.2 Gflops/s (1991). Extensive software development project called Paraphrase is underway. It is focused on program transformations to enable standard FORTRAN programs to run on parallel Cedar machine.

TRAC, (Texas Reconfigurable Array Computer), J.C. Browne et. al., University of Texas, Austin.

16 8-bit microprocessors will be connected via a 4-level banyan switch to 81 memory modules. (JEN81, JEN82, LIP77, PRE82, SEJ80).

CHiP (Configurable Highly Parallel Computer), Lawrence Snyder, Purdue University, Indiana (SNY81/1, SNY81/2, YAL85)

This computer is an array of processing elements embedded in an array of switching elements such that network connectivity between the processing elements can be reconfigured under program control in one machine cycle. The switch lattice is typically a regular structure such as four neighbor or eight neighbor mesh. Fig. 21. illustrates how the original lattice is reconfigured as a mesh and as a binary tree. The project aims to produce 2exp8 and 2exp16 processing elements with a few processing elements on a VLSI chip.

COSMIC CUBE (Nearest Neighbor Concurrent Processor, NNCP), Geoffrey Fox and Charles Seitz, CalTech (California Institute of Technology), Los Angeles, California.

The first machine is 2exp6 Hypercube hosted by VAX11/780, with processor Intel 8086 at each node together with 8087 floating point coprocessor and 128 Kbytes RAM. This machine was commercialized by Intel and marked as iPSC. The Intel iPSC is available with 32, 64 and 128 nodes. Each node is an Intel 80286 processor and 80287 coprocessor together with 512 Kbytes of local memory. The maximal performance of the 2exp10 hypercube is estimated to be about 100 Mflops/s that is to say about the same as the large supercomputers Cray X-MP and Cyber 205. (CHA86, SEI85, EMM85, EMM86/1, EMM86/2).




Fig. 3.: Architecture of Delft Parallel Processor DPP81 (SIP84).


DPP87 (Delft Parallel Processor 87)

This computer is an upgrade of DPP81. The DPP81 consists of one PE-cluster with 8 processing elements (fig. 3).

DPP87 is a modular MIMD system with up to 16 processing modules each having 32 processing elements. Each processing element consists of a stack oriented arithmetic processor AMD 9511. A PDP 11/23 is a host computer. The DPP87 computer is designed for simulation of systems (SIP84).

EGPA (Erlangen General Purpose Array). W. Handler (HAN85)

EGPA consists of a grid-like array of memory-coupled processor modules. Above the array there is a pyramidal hierarchy of processors for supervising and for data transport. Each node consists of one processor and one memory block (fig.4).



- ⊗ Processor-Memory-Module (PMM)
- symmetric multiport-memory connections between neighboring PMMs
- asymmetric multiport-memory connections between PMMs of different hierarchical levels
- ↔ I/O communication to elementary pyramid, supported by I/O processor

Fig. 4.: The EGPA multiprocessor architecture consisting of 85 processor-memory-modules (HAN85).

The project started in 1975. The processor-memory modules are commercially available computers AEG 80/60. Interprocessor communication takes place via common control blocks and mailbox techniques.

FEM (Finite Element Machine), David Loendorf and Harry Jordan, NASA Langley Research Center, Hampton, Virginia. Processing elements are TI9900 microcomputers, controlled by a TI990 minicomputer.

FLEX/32 (Flexible Computer)

This parallel system is composed of 32 bit processors NS 32032. Up to 20 processor modules are connected by a bus and for a box. Common busses link as many as 10 local buses per cabinet (fig 5.).

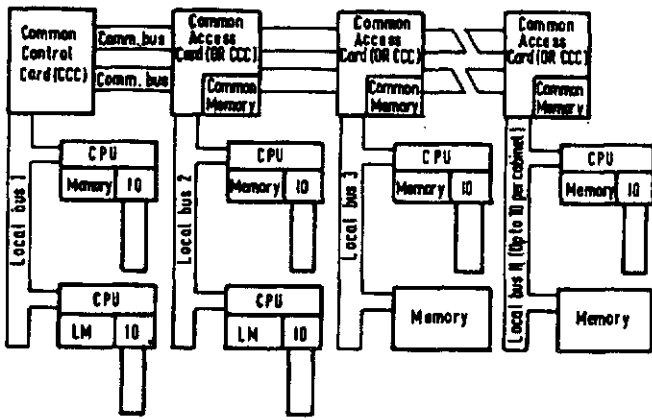


Fig. 5.: The architecture of Flex/32 multicomputer.

A total system is designed for 2480 processors. The system has global and local memory. The price for minimal configuration is \$150,000. (MAN85, Z90B6).

IBM GF11, IBM T.J. Watson Research Center, Yorktown Heights, NY, USA (BEE85) GF11 is a parallel computer with 576 floating point processors (512 primary processors and 64 spares). Each processor has space for 2 Mbyte of memory and is capable of 20 Mflops, giving the total machine a peak of 1.1 Gbyte of memory and 11.5 Gflops. The floating point processors are interconnected by a dynamically reconfigurable non-blocking switching network called the Memphis switch (fig 6.).

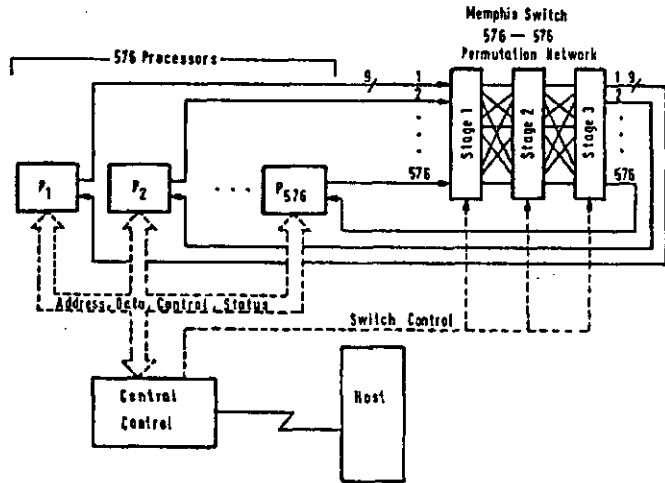


Fig. 6.: The GF11 architecture.

The main intended application of GF11 is a class of calculations arising from quantum chromodynamics in nuclear physics where GF11 is expected to be 100 times faster than Cray 1.

IBM RP3 (Research Parallel Processor Project), IBM T. J. Watson Research Center, Yorktown Heights, N.Y. (PF186).

RP3 project is performed in cooperation with the Courant Institute of Mathematical Science at New York University. The goal of RP3 project is a parallel system with 512 32-bit microprocessors and 2 Gbytes of main storage. RP3 has two multistage Omega-like networks to connect processor-memory elements. The first network is constructed from high-speed bipolar logic and designed for fast interconnection to nonlocal storage. The second network is using a

technique, developed at the NYU Courant Institute Ultracomputer project (EDL85, GOT82, GOT83, SCH80). This network contains more complex functions required to carry out synchronization operation and storage requests. With these techniques the efficiency of parallel system is not degraded as the number of processors is increased.

The language support, initially envisaged, consists of minimal extensions of commonly used languages such as Fortran, C, and Pascal. The adaptation of other languages, such as Common Lisp and Ada to the highly parallel environment is also being studied.

The operating system will probably be an extension of BBD 4.2 Unix, modified internally to make it a fully distributed, symmetrical system, and extended to provide multiple process shared memory and efficient message passing.

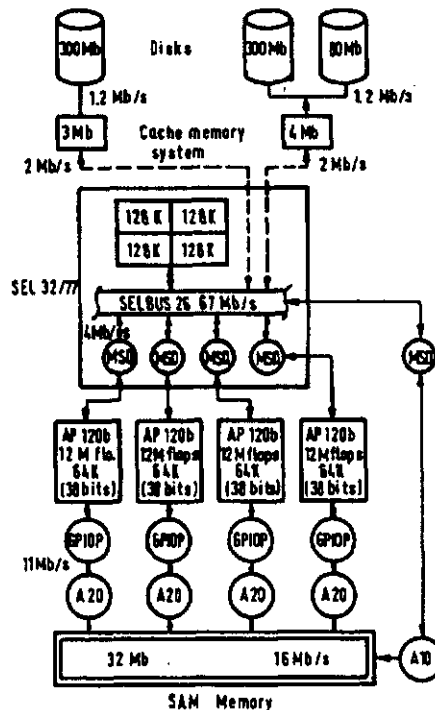
The RP3 is expected to achieve an aggregate of 1000 MIPS on shared memory scientific applications. A demonstration of the RP3 with 64 processor-memory elements is planned for 1987.

RP3 machine is not associated with any IBM product program, and it is expected that very few of the RP3 prototypes will be constructed. RP3 is target to exploit automatic parallelization of Fortran and compilation of functionally-oriented languages.

Development of parallel applications for RP3 is performed before the completion of RP3 hardware construction by means of an experimental emulation system called EPEX (Environment for Parallel Execution).

MIDAS (Modular Interactive Data Analysis System), Creve Maples, University of California, Berkeley.

This is a hierarchical system of a primary



- AP 120B - array processor
- HSD - high speed data interface with 4 Mbits/s throughput
- GPIOP - general purpose I/O processor

Fig. 7.: The architecture of ONERA parallel computer (LECB6).

computer controlling several secondary computers, each of which controls a Multiple Processor Array (MPA). Each MPA also has an input and an output processor and a crossbar switch connecting the processors to 16 switchable memory modules each of 256 Kbytes. The system is used to solve problems in computational physics particularly nuclear science.

ONERA. (Office National d'Etudes et Recherches Aeronautiques).

This is French project on a multi-array processor which started in 1979 and evolved to a loosely coupled architecture. It was designed to solve partial differential equations. The present system has as a host a 32-bit Gould SEL 32/77 minicomputer. Four array processors AP 120B are connected to the SEL bus. The APs are connected to a sharable memory of 32Mbyte (fig. 7.) (LEC86, ADE85)

PASM (Partitioned SIMD/MIMD machine), Howard J. Siegel, Purdue University Indiana. Envisaged that a full machine might have 1024 processing elements connected via a multilevel switch network. A prototype will have 16 Motorola 68000 processors and four control units (SIE81, SCH86).

PRINGLE, University of Washington and University of Purdue (called RP2). 64 processors Intel 8031 are connected via a switch and controlled by an Intel 8086. (KAP84).

SUPERNUM. W. K. Giloi, H. Muhlenbein, (GIL86, HOP86, KR87)

This is a national project in parallel processing in Germany. The computer has a hierarchical hardware structure consisting of nodes, clusters and hyperclusters. It is conceived as a 64 cluster machine with 1024 nodes. Each cluster has 16 nodes and 140 Mbytes winchester disc. Motorola 68020 processors are used in nodes. Fig. 8 shows a SUPERNUM architecture with 16 clusters, where each cluster has 4 processors. The processors in the cluster are connected by a bus. The clusters are connected by row and column rings.




Fig. 8.: The topology of the SUPRNUM parallel computer (KR87).

The system is designed for solving partial differential equations and other numerical applications.

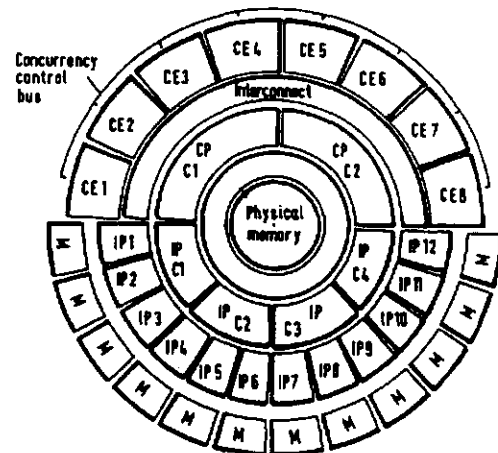
VFPP (Very Fast Parallel Processor), Norman Christ and Anthony Terrano, Columbia University, New York. This is conceived as a 16x16 array of 256 processing elements with nearest neighbor connections. Each processing element has an Intel 80286 with 80287 coprocessor, 20 Kbytes

of local memory and an 11 stage pipelined microprogrammable vector processor. VFPP is a special purpose computer designed for lattice-gauge and similar calculations.

SMALL SCALE PARALLEL SYSTEMS

The word small in this context means only a few processors connected together. To get high processing power by these machines one needs powerful processors with interconnection mechanism which need not be as sophisticated as by connecting several hundred processors.

ALLIANT FX/B (Alliant Computer Systems Corp., Acton, MA 01720) (LLUB6/2, HAR86, SIE86, TEC85) An FX/B computer combines vector and concurrent processing in a system consisting from 1 to 8 computational elements. Each computational element is a microprogramed processor which can execute both scalar and vector instructions. Computational element can access via a crossbar switch, two 64 Kbyte caches. FX/B can have from 2 to 12 interactive processors based on a Motorola 68012 microprocessor. Each interactive processor has 512 Kbyte local memory and is designed for execution of parallel I/O and operating system tasks (fig. 9).



- CE - computational element
- IP - interactive processor
- IP C - IP cache
- CP C - CP cache
- M - multibus

Fig. 9.: The Alliant FX/B architecture.

FX/B operating system Concentrix is an extension of UNIX Berkeley 4.2 version. A minimal configuration with 1 computational element costs \$270,000.

CONVEX C-1 (Convex Computer Corp., Richardson, TX 75081)

C-1 supercomputer is based on a Cray-like architecture. The processing units are interconnected through 64-bit buses and include a dedicated scalar and vector unit. It has a dual ported main memory and up to five 32-bit I/O processors. The operating system is Convex UNIX operating system, similar to UNIX 4.2 BSD operating system. The price for a basic system is \$500,000. New versions are coming to the market. C1 XP processor, a faster version of C-1. (HOL85/3, TEC86)

CRAY X-MP, Steve Chen, Cray Research Inc., Mendota Heights, Minnesota. (HWAB5, OED86, ERH86, LUB85)

The Cray X-MP is a multiprocessor upgrade of the Cray-1 architecture. It comprises 1, 2 or 4 CPUs sharing common memory of up to 8 Mwords in 64 banks of 38 ns ECL memory chips (fig 10).

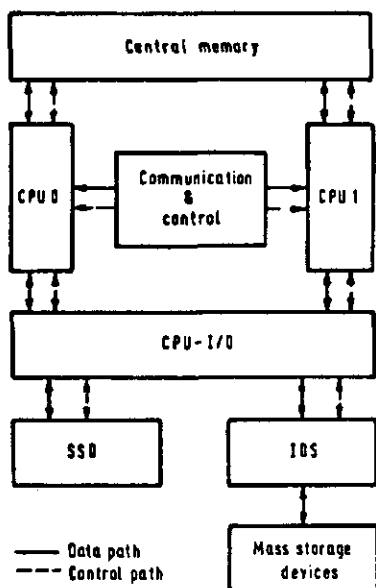


Fig. 10.: Cray X-MP system organization (HWAB5).

Each CPU has 13 pipelined functional units which operate with data in 8 vector registers, each holding up to 64-bit elements. Three memory access pipelines or ports are provided, allowing each CPU to read two vector arguments and store one vector result simultaneously. Each CPU executes its own instruction stream, and rapid synchronization is achieved via 16 shared registers and 32 shared one-bit flags. The clock period is 9.5 ns, giving maximum performance of a pipeline 105 Mflops/s.

CRAY-3, Seymour Cray, Chippewa Falls, Wisconsin. (ERH86, HWAB5, OED86.)

The Cray-3 is scheduled for 1987. It consists of 4 CPUs accessing a shared memory of 256 Mwords. It is an implementation of Cray-2 in gallium arsenide technology, and it is speculated that a clock period of 1ns might be obtained leading to a maximal performance of 1Gflops/s per floating point pipeline.

C.mmp, Carnegie-Mellon University, Pittsburgh, Pennsylvania. (HWAB5, JON80, MAS82, OSL82). This was one of the most ambitious early examples of MIMD computers. This comprised 16 DEC PDP-11 minicomputers connected to 16 memory modules by a 16x16 crossbar switch. The design started in 1971 and the machine was completed in 1975.

Cm*, Carnegie-Mellon University, Pittsburgh, Pennsylvania. This computer was a successor of C.mmp and was based on microprocessors that had now been available. Communication between the microprocessors is via a hierarchical packed switching network. A basic computer module was DEC LSI-11 microprocessor and may act as an independent computer or may be linked to a common interclustered bus with up to 14 other modules to form a tightly coupled cluster. The total Cm* is built-up by loosely coupled clusters. (SWA77).

CYBERPLUS, CDC Corporation, Minneapolis, Minnesota, USA.

The architecture of CyberPlus is based on communication via multiple ring topology. The architecture was derived from the Advanced Flexible Processor which was build for rapid analysis of photographs taken from aircraft. CyberPlus comprises from 1 to 16 CyberPlus processors connected in a ring and attached to a channel of a host CDC Cyber 170/800. Up to 4 such rings can be attached to the host. Communication between processors is achieved by sending information packets to the ring. The packets move round the ring at the rate of one station per clock period, until their destination is reached.

The CyberPlus processor has 256K or 512K words of 64-bit memory for floating point data, 16K words of 16-bit memory for integers and a program memory of 4K words of 240-bit instructions. It has 15 independent functional units. The clock period is 20 ns, giving a floating point capability of 65Mflops/s in 64-mode and 103 Mflops/s in 32-bit mode.

ELXSI 6400, Elxsi Corp., San Jose, California. This computer is similar to CyberPlus. It contains 1 to 12 CPUs and 1 to 4 I/O processors accessing 1 to 6 memory systems via the Gigabus. Potentially the system can achieve 72 Mips. The system incorporates three operating systems: Embos, UNIX BSD 4.3 and UNIX System V.2 - they can all run concurrently. The price for 12-CPU system is approximately \$3 million.

HEP (Heterogeneous Element Processor), Burton J. Smith, Denelcor Inc. Aurora, Colorado, USA. (HWAB5, LING5, SNE85)

The HEP computer was the first commercial computer to offer the facility of programming with multiple instruction streams. A full HEP configuration comprises of 16 Process Execution Modules (PEMs) connected to 128 Data Memory Modules via a multilevel packet switching network called shuttle network. Each may have up to 50 user instruction streams. But the largest system built at the time of writing has 4 PEMs and 4 DMMS and is installed at the NASA Goddard Space Flight Center (fig. 11).

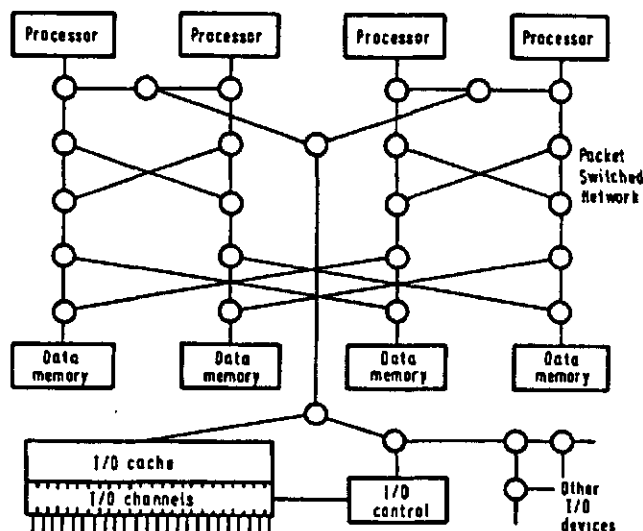
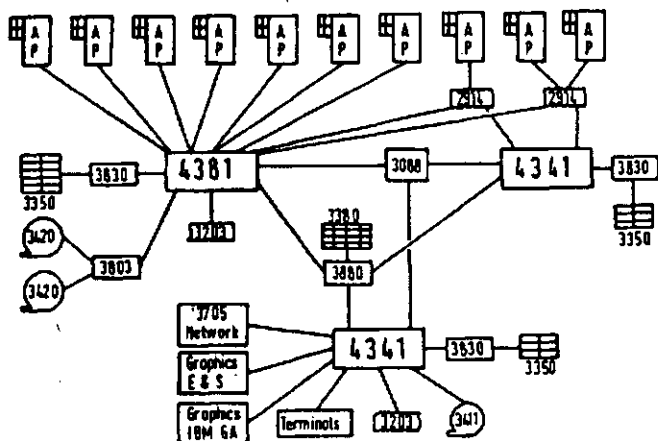


Fig. 11.: The architecture of a typical HEP system with four processors (HWAB5).

IBM LCAP (Loosely Coupled Array of processors), Enrico Clementi, IBM Kingston, USA. (CLE84, NEW86, SIE86) LCAP is a powerful parallel system put together

from parts "off the shelf". These parts are from IBM and Floating Point Systems, Beaverton, Oregon. IBM contributed a host which may be different in different configurations as for example IBM 4381, IBM 4341, IBM 3081, IBM 3089, IBM 3090, and Floating Point Systems contributed attached pipeline processors FPS 164, FPS-164/MAX and FPS-264. Each FPS-* is a single instruction stream computer. An example of LCAP configuration consists of seven FPS-164, each with 4Mbyte of main memory, attached to an IBM 4381 host through a 3 Mbyte/s channel, three FPS-164 are hosted by an IBM-4341 (fig. 12). For favorable problems the system is capable of 60 Mflops/s.



AP - array processor

Fig. 12.: Schematic diagram of the LCAP architecture (CLEB4).

A bottleneck in this architecture is the time needed to transfer information between the constituent computers. It is necessary to decompose a problem into substantial parts that seldom need to communicate with each other. Many physical and chemical problems decompose well and for those problems an LCAP is a cost effective computer.

A similar project to Clementi's is that of Ken Wilson at Cornell University. He has linked 8 FPS-100 to VAX 11/750. He plans to expand the system to 4000 164/MAX to give a theoretical maximum performance of 40 Gflops/s. New versions of LCAP are under construction LCAP-2 and LCAP-3.

MINERVA, Lawrence Widoes, Stanford University, California.

8 Intel 8080 microprocessors and four Intel 3000 microprocessors form a shared memory bus system.

FLURIBUS

This is a symmetric tightly coupled multiprocessor which was based on Lockheed SUE minicomputer, this is a 16 bit computer similar to DEC PDP-11. The system has its beginnings in 1972. A lot of attention was paid to software development. (KAR82)

S-1, Michael Farnwald, George Michael et al., Lawrence Livermore Laboratory, Livermore, California. (HW85)

This is the largest MIMD project, sponsored by US Navy and Department of Energy. The complete design for S-1 computer comprises 16 Cray-1 class pipelined vector computers connected to 16 memory banks by a full cross-bar switch. The S-1 can therefore be regarded as a "grown-up" version of C.mmp. An overall performance is expected to be 1 Gflop/s. Each of the uniprocessors is provided with a data cache of

64 Kbytes and instruction cache of 16 Kbytes in order to limit traffic through the switch. Each memory module may contain up to one Gbyte of storage, giving a total physical storage of up to 2 Gbyte. Single instructions are provided for some common mathematical functions e.g. sine, exponential, etc, and operations e.g. matrix multiply, fast Fourier transform, etc.

We have in this survey excluded logic machines that are being proposed to support the aims of Fifth Generation project, and also dataflow and reduction machines. The reason for this is the fact that these type of computers will probably form a special group of dedicated machines and will not evolve in a general purpose parallel computer of 90'.

3. MULTIPROCESSOR SYSTEM

After we have looked at different parallel computer system architectures, we are going to concentrate on multiprocessors. We are particularly interested in multiprocessors because they are almost general purpose parallel computers and have therefore a great potential power to upgrade or even replace some existing computer architectures.

Multiprocessor is a single computer with multiple processors. Processors communicate and cooperate at different levels in solving a given problem. Multiprocessor is classified as MIMD computer which is defined to be a control-flow computer capable of processing more than one stream of instructions. The communication between processors may occur by sending messages from one processor to the other or by sharing a common memory. Processors have access to common sets of memory modules and peripheral devices.

A multiprocessor system is controlled by one operating system which provides interaction between processors and their programs at the process, data set and data element level.

Multiprocessors are classified according to organizational classification into tightly coupled and loosely coupled multiprocessors and according to structural classification into groups which have similar interconnection structures or topology.

3.1 Multiprocessor organization

Multiprocessors can be organized in a tightly coupled organization or in a loosely coupled organization.

Tightly coupled multiprocessors

Tightly coupled systems can tolerate a high degree of interaction between tasks performed on different processors. Processors communicate through a shared main memory. A small local cache memory may exist in each processor. The connectivity may be accomplished by different interconnection structures between the processors and the shared memory. When two or more processors attempt to access the same memory unit concurrently performance degradation occurs due to memory contention (fig. 13).

3.2. Interconnection structure


Interconnection structure between the memories and processors is either:

- time shared common bus;
- switched network,
 - * crossbar
 - * multistage
 - . omega
 - . banyan
- multiport memories,
- interconnection network,
 - * mesh,
 - * cube,
 - * reconfigurable,
 - * hierarchical.

A term "interconnection network" is sometimes used also for switched networks especially for multistage switched networks because they interconnect processors with memory modules. But we distinguish in this paper switched networks from all other interconnection networks.

Time shared common bus

The time shared common bus is the simplest and attaches every processor to every memory board (fig. 15). Bus requester, driver, and receiver perform all address and data handling. Because its low cost, low complexity and high maximum throughput, the common bus interconnection structure is today the most widely used commercial type of parallel computer system, but is limited to small shared memory computers with up to 20 processor modules (ALLIANT FX/B, ELXSI 6400, FLEX/32, MINERVA).



P - processor
M - memory module

Fig. 15.: A common bus interconnection structure.

Switched network

The switched network system is realized either as a crossbar switch or as a multistage switch.

A further subdivision of switched networks is divided, according to the type of interconnection network, into:

- cross bar (C.mmp, S-1)
- multi-stage (Butterfly, Cedar, GF-11, HEP, RP3, ULTRA, TRAC).

A switched system can be realized as shared memory system or as a distributed memory system. In a shared memory system each processor communicates with each memory module through a switch. In a distributed memory

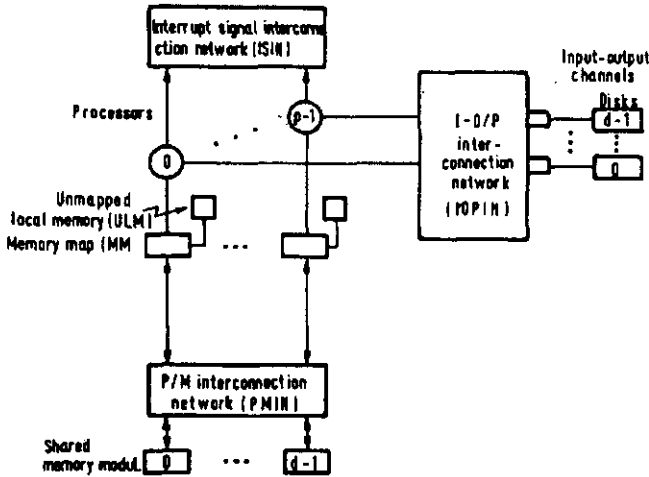


Fig. 13.: Tightly coupled multiprocessors where a complete connectivity exists between the processors and memory (HWAB5).

Loosely coupled multiprocessors

Loosely coupled systems are efficient when the interactions between tasks are minimal. Processes which execute on different computer modules communicate by exchanging messages through a message - transfer system (Fig. 14). In loosely coupled multiprocessors each processor has a set of input-output devices and a large local memory where it accesses most of instructions and data. Sometimes loosely coupled multiprocessors are referred to as a distributed system.

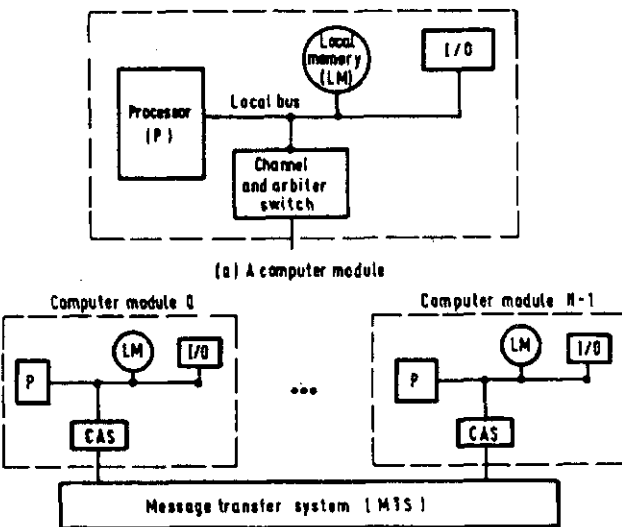


Fig. 14.: Loosely coupled multiprocessors where processes communicate through message transfer system (HWAB5).

The message - transfer system could be a time shared bus or a shared memory system.

Very often computer architecture in parallel processing is a combination of loosely and tightly coupled processors. Loosely coupled systems have often hierarchical organization.

system each memory module is connected as local memory to corresponding processor. The role of a switch is now to interconnect the processing elements, and there are no memory modules connected directly to the switch.

The crossbar switch is an extension of the common bus and implements N buses for N processors and M buses for M memory modules. A separate switch unit connects together a number of processors (P) and memory modules (M). The nodal circuits that couple the processor bus to a memory bus are the switch. A crossbar switch allows all processors to access memory modules simultaneously, as long as each processor accesses different memory module. When two or more processors contend for the same memory module, arbitration lets one processor proceed while the others wait by applying the same techniques as used on a common bus architecture. Since most of the logic is concentrated in the switch nodes, the complexity of a crossbar switch and its cost grows as the square of configuration size. The switch is quite likely to be the largest unit in the system and may be as expensive as one or several of the processors. It is a good choice for systems that are not highly parallel and have about 10 powerful processors.




Fig. 16.: A crossbar switch interconnection structure.

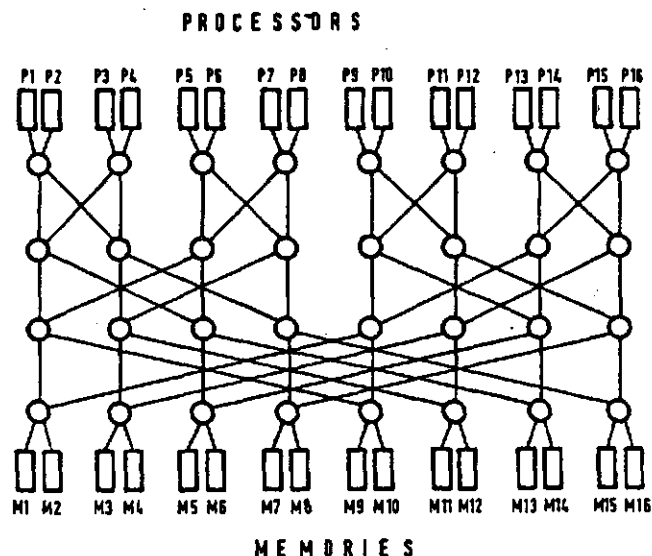
A multistage network reduces the size and the cost of a crossbar switch by linking multiple crossbars as nodes in a network so that each node in a multistage network resembles a small crossbar switch. For example a multistage network that connects 16 processors to 16 memories, realized in two levels consists of 4-by-4 crossbars. The switching elements are distributed throughout the system (fig. 17).

A cost of multistage network that attaches n processors to n memories grows as $n \log n$. All processors can access memory simultaneously, provided that no two processors try to take the same output path from a particular node. In order to reduce this limitation many multistage network architectures have extra pathways to reduce the potential of contention.

The switch in a multistage network is complex. A type of multi-stage switching network is usually omega or banyan. It seems that the multistage shared memory switched computer is the most favored current architecture in parallel processing and enables efficient parallel systems with up to a few hundred processors.

In most switched systems there is both substantial local memory as well as substantial

global memory. Local memory is often realized in a form of registers, cache or buffer memory.



P - processor
M - memory module

Fig. 17.: A multistage network interconnection structure.

Multiport memory

A multiport memory with m ports is similar to n -by- m crossbar switch. In a multiport memory system is the switching logic simply bounded onto the memory module.

Frequently memory module has two ports, one connected directly as local memory to one of the processors, and the other port connected to the switch. Thus each memory bank is both local to one of the processors and globally available to the other processors via the switch.

A multiport memory system could also be treated as a special technical implementation of a switched network and not as a topology.

Interconnection networks

This is a large variety of different interconnection network topologies in parallel systems. Networks are constructed in four different topologies:

- mesh networks (CYBER PLUS, VFPP)
- cube networks (COSMIC CUBE, IPSC)
- reconfigurable network (CHIIP) and
- hierarchical network (Cm*, EGPA, SUPERNUM).

Mesh networks are one or multidimensional and are realized in square, hexagonal or other geometry. In a square mesh of dimensionality, d , each processing element is connected to $2d$ neighbors. The number of processing elements $N = n \exp d$, may be varied independent of the dimensionality by increasing the linear dimension of mesh n . For example, a square mesh with dimensionality 2 is connected to four neighbors. Fig. 18. shows a mesh system with 16 processing elements.




Fig. 22.: Linear speedup in a multistage switched network with 256 processors for Butterfly parallel system.

4. APPLICATION OF PARALLEL SYSTEMS

Uniprocessor architectures are approaching theoretical limits in processing speed.

In high speed or real time processing tightly coupled computer systems have to be used.

Most parallel computer nowadays are designed for numerical work with floating point numbers, and are build for the solution of large problems in physics, chemistry and engineering.

Large computer capabilities are necessary particularly in:

- complex graphic images,
- structural analysis,
- aerodynamics,
- meteorology,
- medical diagnostics,
- research in an oil exploration,
- research in fusion physics,
- industrial automatization,
- processing of sensing signals,
- genetic engineering,
- molecular dynamics,
- quantum mechanical problems,
- socioeconomic models, etc.

Mathematical problems which are solved by parallel systems are:

- Monte Carlo simulation,
- Hartree-Fock equation in the electron gas,
- finite element methods, etc.

Many of multiprocessors are almost general purpose as for example:

ALLIANT, BUTTERFLY, CEDAR, C.mmp, Cm*, CONVEX, COSMIC CUBE, CRAY X-MP, CRAY-3, CYBERPLUS, DCA, DPP, EGPA, ELXSI 6400, FLEX/32, FMP, HEP, IBM RP3, IBM LCAP, IBM GF11, MINERVA, ONERA, PLURIBUS, PRINGLE, SUPERNUM, S-1, TRAC, ULTRA.

These multiprocessors are designed for large scientific and engineering problems, for CAD automation, real time voice data multiplexing and other computationally involved problems.

Some multiprocessors are more limited in applications and are considered as special purpose parallel computers designed for one bit logic operations, or image processing, knowledge based expert systems or designed for other special applications in artificial intelligence. Special purpose multiprocessors are:

CHIP, DADO, FEM, MANIP, MEIKO, PASM, PUMPS, VFPP.

5. WILL PARALLEL PROCESSING WIN?

Many ambitious projects in parallel processing in past have failed. For example ILLIAC IV cost four times the original contract figure and did not come even within a factor of 10 of its originally proposed performance. However its influence was profound and ILLIAC IV was the first to pioneer the new and faster emitter-coupled logic (ECL) rather than the established transistor-transistor logic (TTL). ILLIAC IV also pioneered the use of 15-layer circuit boards and computer aided layout methods.

Other parallel computer systems of the 70' were also not very successful. For example C.mmp and Cm* had problems with hot memories because their interconnections structure, which was based on crossbar switch was not intelligent and could not overcome this problem. Nowadays these solutions are given. BSP NASF had a bottleneck in a central control processor and no efficient synchronization mechanisms were known at that time. But again ICL DAP was pioneering in an important feature of engineering design that processing element logic is mounted on the same printed circuit board as the memory to which it belongs. VLSI technology can now include processing element and its memory on the same chip.

Now the technology has advanced sufficiently to make parallel architecture practicable. Therefore we see such great interest in parallel processing.

6. CONCLUSION

It is not possible to predict which of these varied computer architectures will prove the most successful in future on the market.

By analyzing the performance of multiprocessors which is primarily dependent on interconnection structure one might get an insight to the development of parallel computer systems and try to predict future trends in parallel computing. It seems that in next decade the most influence on parallel computing are going to have the projects in massively parallel processor:

- NYU Ultracomputer, whose principles are applied in IBM-RP3 parallel system,
- Butterfly, produced by the company BBN Bolt, Beranek & Newman,
- Cedar, a multiprocessor supercomputer of the University of Illinois.




Fig. 18.: A mesh network with 16 processing element connected in a lattice (a) and connected as a torus (b).

Cube networks have either hypercube architecture or cube-connected-cycles network architecture. A cube-connected-cycles network is a cube where each node of the hypercube is replaced by a ring (or cycle) of processing elements. In a d-dimensional binary hypercube there are d connections to each processing element, $n=2$, and therefore the number of processors equals $N=2 \exp d$. We see that the number of processing elements can not be increased without also increasing the number of connections to each processing element. For example, a six-dimensional hypercube which has 64 nodes is topologically the same as $4 \times 4 \times 4$ three dimensional mesh with triply periodic boundary conditions.




Fig. 19.: A cube network with 4, 8 and 16 processing elements.

Hierarchical class of multiprocessor systems is realized as tree network, hierarchy of pyramids or clusters of clusters.






Fig. 20.: A hierarchical network realized as a tree network.

Reconfigurable networks include all cases in which the interconnection pattern between processing elements can be changed. This is usually achieved by interspersing switching elements between the processing elements which may be controlled by a user program.

the original lattice architecture



the switch lattice configured as a mesh



the switch lattice configured as a binary tree




Fig. 21.: The original switch lattice in CHIP parallel computer configured as a mesh and as a binary tree.

Let us compare three most widely used topologies: common bus, crossbar and multistage network. Seven features are going to be compared:

- 1 - cost
- 2 - complexity
- 3 - max. throughput
- 4 - interconnect bandwidth
- 5 - # of signal paths
- 6 - efficiency
- 7 - max. # of CPU

We see (table 1.) that cost of a parallel system is the lowest in a common bus topology, but efficiency drops with increasing the number of processors. A crossbar switch is very powerful in connecting a few processors, but the price and complexity of a system is very high. A multistage network is a good topology to interconnect a large number of processors for a medium cost.

FEATURE	BUS	CROSSBAR	MULTISTAGE NET.
1	low	high (CPU exp ²)	medium (n logn)
2	low	high (CPU exp ²)	medium
3	high	no limit	high
4	fixed by cycle time	proportional to # of CPU	proportional to # of CPU
5	large	medium	medium
6	drops	linear	linear
7	up to 30	up to 10	up to 1000

Table 1.: Comparison of different features for a common bus, a crossbar and a multistage parallel computer system.

New synchronization mechanisms for multistage switched networks are nowadays enabling almost linear speedup for systems with up to 256 processors (fig. 22).

The reasons for success of these projects seem to be the fact that they are devoted to development of a GENERAL PURPOSE MIMD parallel computer. Excellent performance results are reached particularly because they use:

- MULTISTAGE INTERCONNECTION NETWORK: A near linear speedup is reached by 256 processors using as interconnection structure between the memories and processor a multistage interconnection network (Omega network);

- INTERLEAVING: that is spread data uniformly throughout common memory modules in order to avoid contention for any one memory module;

- FETCH-AND-ADD: a very effective interprocessor synchronization operation.

An operating system seem to be a parallel version of a UNIX-like operating system. It is possible to achieve high performance by connection a large number of processing elements, even with "of the shelf" standard processors. It seems also that some architectural features as for example the size of local memory or the size of cache memory at every processor is of secondary importance for high performance of a parallel system.

7. REFERENCES

- (ABU84) Abu-Sufah W., A. Kwok, Performance Prediction Tools for Cedar: a Multiprocessor Supercomputer, IEEE Conf. on Comp. Architecture, 1984, p.406-413
- (ABU86) Abu-Sufah W., H. Husmann, D. Kuck, On I/O Speedup in Tightly Coupled Multiprocessor, IEEE Trans. on Computers, June 1986, p. 520-530
- (ADE85) Adelatado M., D. Comte, P. Siron, Ph. Berger, Expression of Concurrency and Parallelism in an MIMD environment, Computer Physics Commentars 37, 1985, p. 63-67, North Holland
- (BAR68) Barnes G. at al., The Illiac IV Computer IEEE Trans. on Comp., August 68, p. 746-756
- (BAT77) Batcher K., The multidimensional Access Memory in STARAN, IEEE Trans. on Comp., 1977, p. 174-177
- (BAT80) Batcher K. E., Design of a Massively Parallel Processor, IEEE Trans on Comp., Sept. 80, p. 836-844
- (BAT82) Batcher K. E., Bit Serial Parallel Processing System, IEEE Trans. on Comp., May 82, p. 377-384
- (BEE85) Beeten John, et al., The GF11 Supercomputer, IEEE, pp 108-115, 1985.
- (BOU72) Bouknight at al., The Illiac IV System Proc. IEEE, April 1972, p. 369-388
- (CLE84) Clementi E. at al. Parallelism in Computations in Quantum and Statistical Mechanics, Proceeding on 2nd International Conf. on Vector and Parallel processors in Comp. Sci., Oxford, August 84, p.287-294
- (CHA86) Chamberlain Richard, Experiences with the Intel iPSC hypercube, Supercomputer, p 24-29, 1986.
- (DAV69) Davis R., The Illiac IV Processing Element, IEEE Trans on Comp., Sept. 69, p. 800-816
- (EDL85) Edler J., A. Gottlieb at.al, Issues Related to MIMD Shared-memory Computers: the NYU Ultracomputer Approach, IEEE Conf. on Comp. Architecture, 1985, p. 126-135
- (EMM85) Emmen ad, Intel's iPSC: a family of parallel computers based on microprocessors, SUPERCOMPUTER News, May 1985
- (EMM86/1) Emmen Ad, Hypercube-toy or tool?, SUPERCOMPUTER News, July/September 1986
- (EMM86/2) Emmen Ad, Vector extension for the iPSC, SUPERCOMPUTER News, July/September 1986
- (ERH86) Erhel J., Parallel programming and applications on Cray X-MP, Supercomputer, Sept. 86, p. 53-60
- (FIN77) Finnila, Charles A., H. Love, The Associative Linear Array Processor, IEEE Trans. on Comp., Feb. 77, p. 112-129
- (GIL86) Giloi W. K., H. Muhlenbeim, Rationale and Concepts on the Supernum Supercomputer Architecture, MIPRO 86, Opatija, 1st Yugoslav Conf. on New Generation of Computers, p. 3.1-3.17
- (GOT82) Gottlieb A. at al., The NYU Ultracomputer - Designing a MIMD Shared Memory Parallel Computer, IEEE Conf. on Comp. Architecture, 1982, p. 27-42
- (GOT83) Gottlieb A. at al., The NYU ULTRA computer-Designing on MIMD shared Memory parallel Computer, IEEE Trans. on Comp., 1983 .
- (HAN85) Handler W. at al., A tightly coupled and hierarchial Multiprocessor architecture, Computer Physics Comm 37, 1985, p. 87-93
- (HAR86) Hars N., New Systems offer near-supercomputer performance, IEEE, March 86, p. 104-107
- (HOC81) Hockney R.W. and C. R. Jesshope, Parallel Computers, Adam Hilger Ltd, Bristol, p. 126-143, 1981
- (HOL85/1) Hollenberg Jaap The Cray-2 computer system SUPERCOMPUTER 8/9 , September 1985
- (HOL85/2) Hollenberg J., The Butterfly Parallel Processor Computer System, Supercomputer, Sept. 85, p. 23-27
- (HOL85/3) Hollenberg J., The C-1: A Minisuper Supercomputer, March 85, p. 7-8
- (HOP86) Hoppe H. C., H. Muhlenbeim, Parallel adaptive full-multigrad methods on message-based multiprocessor, Parallel Computing, Oct. 86, p. 269-289
- (HWAB5) Hwang K. and F. Briggs Computer Architecture and Parallel Processing, McGraw-Hill Book Company, p. 237-241, 1985.
- (JEN81) Jenevein R., D. Degroot, G. Lipovski, A Hardware Support Mechanism for Scheduling Resources in Parallel Machine Environment, IEEE Conf. on Comp. Architecture, 1981, p. 57-65
- (JEN82) Jenevein R., J. Brown, A Control Processor for a Reconfigurable Array Computer, IEEE Conf. on Comp. Architecture, 1982, p.81-89
- (JON80) Jones A., P Schwarz, Experience Using Multiprocessor Systems: A Status Report, ACM Computing Surveys, June 80, p.121-167
- (JOR82) Jordan T. L. A Guide to Parallel Computation and Some Cray-1 Experiences Parallel

Computations AP, 1982

(KAP84) Kapauan A., J. Field, D. Gannon, L. Snyder, The PRINGLE Parallel Computer, IEEE Conf. on Comp. Architecture, 1984, p. 12-20

(KAR82) Kartashev S., S. Kartashev, Designing and programming modern Computers and Systems, vol. 1, chapter II, Prentice-Hall, p. 143-154, 1982.

(KOC85) Koch Wilhelm, First European installation of Siemens VP-200, SUPERCOMPUTER 7, May 1985.

(KOG81) Kogge Peter M. The Architecture of Pipelined Computers, McGraw-Hill Book Company, p. 159-162, 1981.

(KRAB7) Kramer O. and Muhlenbein H., Mapping Strategies in Message Based Multiprocessor Systems (to be published).

(KUC82) Kuck David J. and Richard A. Stokes The Burroughs Scientific Processor (BSP) IEEE TRANSACTIONS ON COMPUTERS vol. C-31, No. 5, May 1982

(LEC86) Leca P., The ONERA experimental MIMD system, Supercomputer, Sept. 86, p.91-96

(LIN82) Lincoln Neil R. Technology and Design Tradeoffs in the Creation of a Modern Supercomputer IEEE TRANSACTIONS ON COMPUTERS vol. C-31, No. 5, May 1982

(LIN85) Linebock R., Parallel Processing: Why a Shakeout News, Electronics, Oct. 85, p. 32-34

(LIP77) Lipovski J., On a Varistructured Array of Microprocessors, IEEE Trans. on Computers, Feb. 1977, p 125-138

(LLU84) Llubra Rossend, VP-200: Fujitsu's Supercomputer, SUPERCOMPUTER 2, July 1984

(LLU86) Llubra R., The Alliant FX/8 entry level supercomputer, SUPERCOMPUTER, March 86, p. 7-11

(MAN85) Manuel Tom, Parallel Machine Expands Indefinitely, Electronics Week, May 85, p. 49-53

(MAS82) Mashburn Henry, The C.mmp/Hydra project: An Architectural Overview, Computer Structures: Reading and Examples, ed D. Sieworek, Bell, Newell, p. 350 - 370, McGraw Hill, 1982

(OED86) Ded W., O. Lang, Modeling, measurements and simulation of memory interference in the Cray X-MP, Parallel Computing, Oct 86, 343-359

(OSL82) Oslund B., P. Hibbard, R. Whiteside, A Case Study in the Application of the Tightly Coupled Multiprocessor to Scientific Computation, Parallel Computations, ed G. Rodrigue, p. 315-364, Academic Press, 1982

(PF186) Pfister G. F., Parallel processor project to link 512 32-bit micros, IEEE Computer, Jan. 86, p 98-99

(PRE82) Premkumar U., J. Browne, Resource Allocation in Rectangular SW Banyans, IEEE Conf. on Comp. Architecture, 1982, p. 326-333

(PUR74) Pursell Charles J. The control data STAR-100- Performance measurements NCC 74

(RET86) Rettberg Randall and Robert Thomas, Contention is no obstacle to shared-memory multiprocessing, Comm. of the ACM, vol. 29, No. 12 p.1202-1212, December 1986.

(RUD72) Rudolph J., A production implementation of an associative array processor- STARAN, Fall Joint Computer Conference, 1972, p. 229-241

(RUS78) Russell Richard M. The CRAY-1 Computer System Comm. ACM, vol. 21, No. 1, January 1978

(SCH80) Schwartz T., Ultracomputer, ACM Trans. on Programming Languages and Systems, Oct. 1980, p. 484-521

(SCH86) Schwederski Thomas and Siegel Howard Jay, Adaptable Software for Supercomputers, IEEE Computers, pp.40-48, February 1986.

(SEJ80) Sejnowski et al., Overview of the Texas Reconfigurable Array Computer, AFIPS National Computer Conference, 1980, p. 631-642

(SEI85) Seitz Charles L., The Cosmic Cube, Communications of the ACM vol.28, No.1, January 1985

(SIE81) Siegel H., PASM: A Partitionable SIMD/MIMD System for Image Processing and Pattern Recognition, IEEE Trans on Computers, Dec. 1981, p. 934-947

(SIE86) Sieworek D., New Trends in Comp. Architecture, MIPRO Conference, Opatija 1986

(SIP84) Sips H., The DPP81- an exercise in parallel processing, Supercomputer, Nov. 84, p. 31-37

(SNE85) Snelling D., HEP Applications: real time flight Simulation, Computer Physics Comm. 37, 1985, p.261-271

(SNY81/1) Snyder L., Programming Processor Interconnection Structures, Technical Report CDS-TR-381, Perdue University, 1981

(SNY81/2) Snyder L., Introduction to the Configurable, Highly Parallel Computer, IEEE Computer, Jan. 1981, p. 47-56

(SWA77) Swan, Fuller, Sieworek, Cm* - A Modular Multi-microprocessor, AFIPS National Computer Conference, 1977, p637-644

(UCH85) Uchida Keiichiro and Mikio Itoh, High Speed Vector Processor in Japan, Computer Physics Communications 37 (1985) 7-13, NH Amsterdam

(VON84) Vons Peter Cyber 205 vector-features used by vectorizers SUPERCOMPUTER 3, September 1984

(WIL82) Wilson Kenneth G. Experiences with a Floating Point Systems Array Processor Parallel Computation AP, 1982

(YAW77) Yaw S. S., H. S. Fung, Associative Processor Architecture- A Survey, ACM Computing Surveys, March 77, p. 3-27

(ZSO86) Zsohar Leslie et all., Bus Hierarchy Facilitates Parallel Processing in 32-bit Multicomputer, Computer Technology Review, summer 1986, pp 51-59.

(NEW85) News, China's first supercomputer, SUPERCOMPUTER 6, March 1985

(NEW86) News, "Supercomputer is just an advertisement word" - an interview with E. Clementi, SUPERCOMPUTER, Sept. 86, p. 24-33

(TEC86) Technology to watch, This Minisuper is aimed at parallel processing, Electronics, Oct. 86, p. 56-60