ETA SYSTEMS

Piper Computer

System Overview

April 16, 1986

Introduction

The Piper system will consist of a one or two processor computer system, with both processors having access to a shared memory. The computer will directly connect to high performance supercomputer disk and tape subsystems, while relying on network interfaces to a variety of computer systems, terminals, and/or workstations to perform low speed I/O, job introduction, output generation, and support for non-standard peripheral equipment. All peripheral and network connections are through Input/Output Units (IOUs) into shared memory.

The system will include a service unit which provides for operator display and control, system reconfiguration, maintenance functions, and will have removable media for software installation.

Following is a diagram showing the functional components and their interconnection paths. Path widths in bits and bandpass in megabytes (MB) or megabits (mb) is also shown.

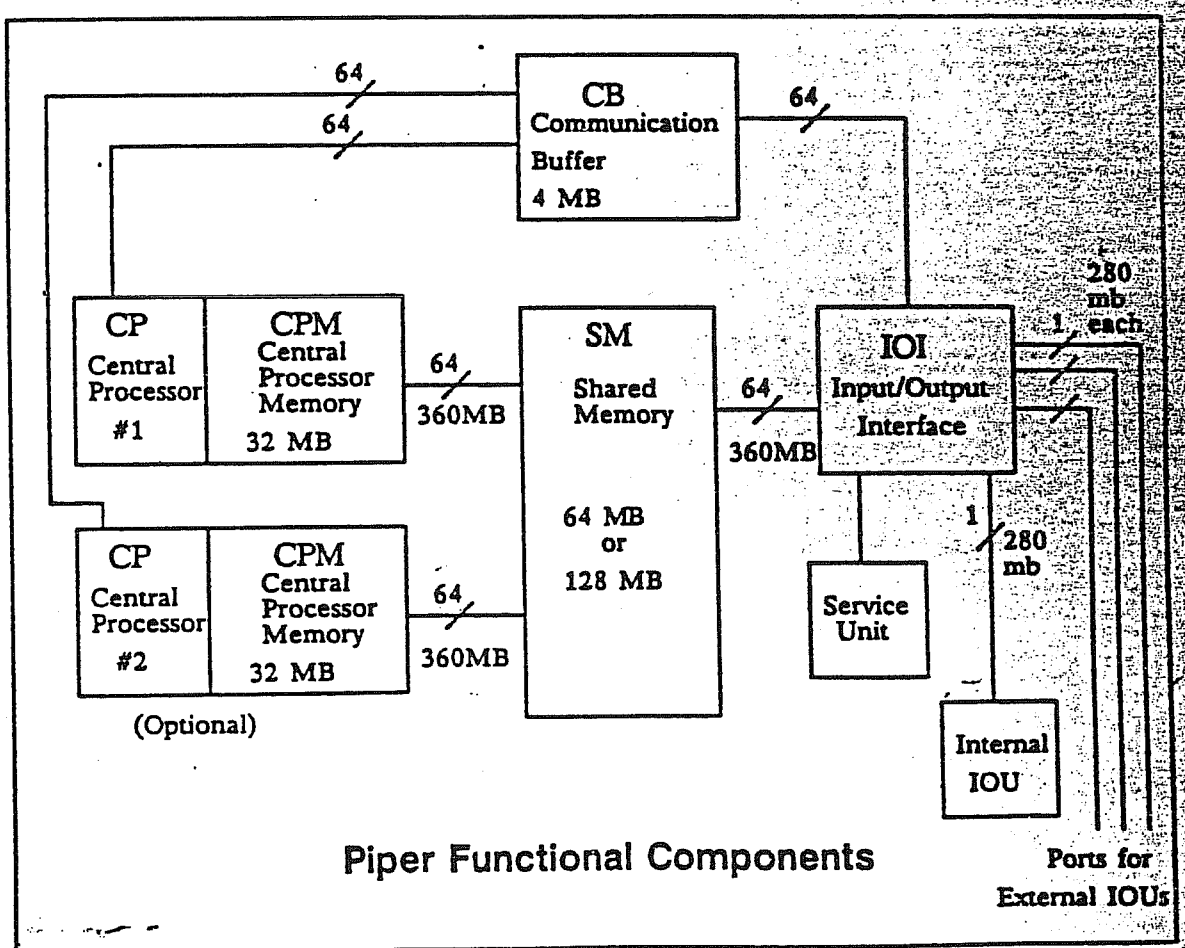

System Components

Below is a brief summary of each Piper subsystem and expected performance data.

Central Processing Unit (CPU)

The Central Processing Unit consists of an ETA^{10} compatible processing unit containing scalar and vector processing capability and virtual memory management hardware. There are 32 megabytes of high speed memory for program and data storage. The processor machine cycle is expected to be about 22 nanoseconds. Processor to memory bandwidth is 64 bytes per machine cycle in 4 separate streams or 2909 megabyes/sec. The scalar processor has a 256 word register file that supplies operands to a scalar arithmetic/logical unit. Scalar results are produced one per machine cycle (46 megaflops peak rate). The vector processor has two arithmetic/logical pipelines that can compute memory to memory results at the rate of 2 per machine cycle in full precision or 64 bit mode (91 megaflops peak) or 4 per machine cycle in half precision or 32 bit mode (182 megaflops peak). Vector element scatter or gather and linked triad (scalar & vector & vector) operations are available in both modes. A half precision linked triad operation can produce results at a 364 megaflop peak rate. A dual processor Piper system can thus produce a peak rate of over 700 megaflops.

Complementary Metal Oxide Semiconductor (CMOS) gate array technology is used in the processor with each gate array containing about 18,000 useable gates. A comprehensive built in self test capability is included on all arrays to provide complete physical defect and chip to chip interconnect testing before and after system assembly. The entire processor is contained on a single printed circuit board. This combination of high density low power circuits bonded onto a single circuit board assures excellent quality and reliability.

System Memory (SyM)

System Memory consists of a Shared Memory and a Communications Buffer and all required interface logic on a unified backpanel.

Shared Memory (SM)

The primary role of Shared Memory will be to provide a high capacity storage unit for application problem storage. Shared Memory sizes of 64 or 128 megabytes are available. There are 3 high speed ports into or out of the memory for allocation to the two processor and the I/O subsystem. Each Shared memory port has a transfer rate of 8 bytes per cycle for a total Shared memory bandwidth of 1090 megabytes/sec.

Communication Buffer (CB)

The Communication Buffer will provide a low overhead mechanism for interprocessor communication. Processes executing on the central processors or in the IOUs will signal each other by way of the Communication Buffer. It will have a fixed size of 4 megabytes. The CP or I/O to CBM bandwidth will be one 8 byte word every machine cycle or 363 megabytes.

Input/Output Subsystem

Input/Output Units (IOUs) will provide for the attachment of perpherials to the Piper computer system. Data through the I/O subsystem port into Shared Memory is controlled by the Input/Output Interface for distribution to the IOUs and the Service Unit. One IOU will be built into the Piper mainframe cabinet. Up to 3 external IOUs, identical to those used in ETA10 computer systems, can be connected via fiber optic cables to the Input/Output Interface. Various functional units that present a specific external interface for peripheral connection can be added to the base IOU.

The serial Data Pipe from IOI to each IOU will operate at 280 megabits, full duplex, with an effective data transfer rate (each direction) of 220 megabits per second at 100 feet (30 meters), reduced as distance increases. Maximum Data Pipe length is 1000 feet (305 meters).

Peripheral Connections

The internal IOU will provide access to its Data Pipe for a set of 5 functional units. These will be identical to those used in the ETA10. The functional units will be the "channel" interface to the I/O subsystem from peripheral devices. One unit will manage the mass storage file system. Others will be available to interface to ISI protocol disks, FIPS-60 protocol devices, Ethernet or the CDC Loosely Coupled Network.

Initial mass storage will be provided by Ibis or CDC Hydra disks connected via a Disk Channel Controller. These will provide a capacity of 1 to 8 1250 or 625 megabyte disks with a 12 megabyte bandwidth on each Disk Channel Controller.

A reel to reel tape capability for mass storage backup or data interchange to other computer systems will be provided. This will be a full 200 inches per second start-stop FIPS-60 compatible tape system.

Networks

The ETA Open Interconnect Network is based upon the IEEE 802.3 standard for Ethernet which has a raw transfer rate of 10 megabits/sec and and effective application transfer rate of 1 to 1.5 megabits/sec. The U. S. Department of Defense Standard TCP/IP protocol is supported over the Open Interconnection Network. Standard application protocols supported include File Transfer Protocol (FTP) and TELNET. The Open Interconnection Network is also designed to be compatible with other vendors' offerings that rely on Ethernet and TCP/IP. This design minimizes the constraints on customers in their selection of workstations or other equipment to be used in conjunction with the Piper computer. As the International Standards Organization specifications for all seven layers of their Open System Interconnection model are implemented in the commercial marketplace, these standards also will be supported.

The Piper computer system will also offer interfaces to the Control Data Loosely Coupled Network (LCN) and Remote Host Facility (RHF) applications. Network access is or will be available for connection to the following computer systems:

CDC Cyber 170/180 with NOS 2.
CDC Cyber 170/180 with NOS/VE.
CDC Cyber 205 with VSOS.
IBM 303X and 43XX with MVS.
Digital Equipment VAX with VMS.

Service Unit (SU)

Each CPU, the Shared Memory Interface, Input/Output Interface, and the IOUs will have a Maintenance Interface to facilitate problem isolation and prediction by the Service Unit.

The Service Unit will consist of a 2 node Apollo network with one utility server node and one color keyboard/display operator's node. These two nodes communicate with each other on the 12-Megabit DOMAIN network. The operators node will have integral disk storage for Apollo system software, and there will be another 500 megabytes of disk storage on the server processor for ETA system software and other storage.

Performance and Price

Overall performance estimates for Piper compared to other machines is shown in the table below. These estimates assume that automatic vectorization software (within compilers or as front-end preprocessors) has been used. The performance factors are for single processor configurations of all machines except the DEC 8800. The DEC machines have no vector capability. The performance factors used are presented as ranges. The lower number represents an estimate for the scalar performance factor, the higher number represents an estimate for the performance factor on typical vectorizable scientific problems. It should be noted that Piper's architecture provides additional performance benefits that are not reflected in the table. The existence of the 8 or 16 million word Shared Memory and the Communications Buffer coupled with the balance between CPU speed, memory bandwidth and I/O bandwidth enhances Piper's ability to run very large problems as well as make Piper a very effective throughput machine for a large class of small problems.

Piper's Overall Performance Compared to Other Machines

Scalar		Vector *		Vector **	
.8	to	1	to	1	times the CDC Cyber 205
.8	to	2.5	to	5	times the CDC Cyber 990
.4	to	.8	to	1.6	times the Cray XMP1
.8	to	1.3	to	2.6	times the IBM 3090-200 VP
1.7	to	3.5	to	7	times the Scientific Computing Systems SCS-40
4	to	7	to	14	times the Convex C1
3	to	100	to	200	times the Digital VAX 8800
20	to	1000	to	2000	times the Digital VAX 11/780 (8200)

* 64 bit mode ** 32 bit mode

The projected price for 1987 installations will be in the range of $2 million to $2.5 million.

Environmental and Power Requirements

The Piper computer system requires 220 volt 30 amp service and a temperature and humidity controlled environment. It can be installed without raised floors and can be wheeled through ordinary doors. It will be capable of meeting accepted noise and radiation emission standards

for typical engineering work areas and of satisfying all United States, Canadian or European environmental and safety standards.

Operating System and Product Set

The Piper Prototype system will execute the ETA^{10} operating system. This virtual memory system will support direct interactive or batch processing. To help programmers create applications for the Piper, the system provides two different programming environments: UNIX, a highly portable and widely used environment; and VSOS, the user interface of the CYBER 205. Each of these environments provides the command language, tool sets, and the most important programming languages to which current users of these standard systems are accustomed. Both of these user environments and any future ones will have access to a set of ETA Systems products that are environment independent including support libraries, symbolic debuggers and multitasking aids.

The UNIX programming environment on the Piper is based on AT&T System V, Release 2, with some Berkeley extensions. It has both the Bourne and C shells. Almost all the utilities and routines included in UNIX System V are supported, including the standard UNIX kernel calls. Additional networking facilities and accounting information also are available for users who wish to access them. Products common to all environments, including the numeric and multitasking libraries, also are available for use with the UNIX programming environment.

The VSOS programming environment is based on the CYBER 205 VSOS Release 2.2. It has been enhanced in several ways. It provides richer interactive features, such as improved job control language with procedure file capabilities and multiple parallel sessions per user. The user interface also has been simplified by removing some CYBER 205 paging file limits and other file system limitations.

Several ETA Systems software products are designed to be used with either programming environment. For this reason, they are called common product languages or common product tools. These products include the standard programming languages (ETA FORTRAN, C, and in the future, Pascal), multitasking capabilities, and some program development tools (debugging tools, for instance).

The overall guiding design objective for ETA FORTRAN is to produce a compiler that makes it easy for programmers to gain access and to take advantage of *every* feature of Piper's architecture. This completely new ETA FORTRAN compiler supports the ANSI 78 standard language and the anticipated array notation of the next ANSI standard. The compiler also supports some of the most popular language extensions from other major vendors, including BUFFER IN/BUFFER OUT, CYBER 205 FORTRAN vector syntax, IBM-compatible arithmetic type statements, and multiple assignment statements. An increase in the performance of existing FORTRAN software will be facilitated by the new ETA FORTRAN along with state of the art automatic vectorization. There are two ETA FORTRAN compilation modes on the Piper: production and development. The production mode stresses execution speed and provides a variety of optimizations and vectorization. The development mode emphasizes compilation speed and extensive diagnostics. Users may select the compilation mode most appropriate for their needs.

The C language, as defined in *The C Programming Language* by Brian Kernighan and Dennis Ritchie, is fully implemented on the Piper system. The C compiler emphasizes compatibility with other portable C compilers, fast compilation speed, and diagnostics.

Hardware and Software System Support

ETA Systems intends to support Piper in the same way that the ETA^{10} will be supported.

The maintenance strategy is built on the reliability of the components of the ETA^{10}. Very few field replaceable units will exist. A remote support center will use Piper's built in facilities to isolate problems. On call repairs can then be made in the field.

All early Piper systems will be provided with analyst support. The long range goal is to provide no on site analyst support; only central analyst support. This is the same strategy being pursued for the ETA^{10}.

Possible Enhancements

The present design of the Piper computer system has included the potential for various future upgrades and enhancements. Some of these are listed below:

Environmental

The Piper computer system mainframe is designed to operate in a typical engineering work area. If peripherals compatible with an office were used, the entire system could be moved to an air conditioned office where 220 volt 30 amp service was available.

Software

A full implementation of Pascal, as defined in the *Pascal User Manual and Report* by Kathleen Jensen and Niklaus Wirth, could be available after the first release of the ETA^{10} system. The language would be extended to include a vector notation. The Pascal compiler would emphasize fast compilation speed and good diagnostics.

Peripherals

New disk equipment can later replace the present more expensive high performance drives. This would likely be up to 4 SMD-E protocol, 800 megabyte capacity, 3 megabyte transfer rate disks attached to a VME Channel Controller. An alternative might be to support a FIPS-60 compatible disk systems such as the IBM 3380 system.

A low cost medium performance streaming tape drive with modest start-stop capability could be included. It would interface to a VME Channel Controller via the industry standard Pertec or STC-1935 tape interface protocols.

The possibility also exists to add any custom VMEbus based controller and appropriate driver software to the IOU and the operating system kernal to connect other special equipments. It is possible to add NSC Hyperchannel network access using such a VMEbus controller and NETEX applications.

Networks

In addition to directly tapping an Ethernet Network, ETA Systems will offer a variety of devices that facilitate the connection of workstations, computers, terminals, and output devices to the network. These products could include:

Protocol processing printed circuit boards and generic software for insertion into computers and workstations having an open architecture and a common bus interface, such as Multibus, Unibus, Q-bus, VME bus, or PC/AT bus.

Terminal servers offering from 4 to 32 RS232, RS422, or V.35 ports for the connection of terminals, or for passthrough into other host mainframes supporting those terminals.

Gateway servers supporting an X.25 interface, either to a public data network or to a host computer system supporting X.25 protocols.

Bridges for connecting multiple Ethernet segments.

Print stations and spooling software so print facilities are available directly from the Piper.

Network maintenance stations that provide control of network configuration, usage recording, directory services, and diagnostics.

The following diagram depicts an example of a Piper system configuration with a variety of devices connected by the Open Interconnection Network.