

# ETA SYSTEMS

# **ETA SYSTEMS**

**- DESIGN -**

**- DEVELOP -**

**- MANUFACTURE -**

**- MARKET -**

**- SUPPORT -**

**The World's Fastest and Most  
Reliable Supercomputer System**

# CLASS VI & CLASS VII SUPERCOMPUTERS

## Class VI: Current Generation

- “Class VI” – Dept. of Energy term
- Single CPU system with vector processing
- Memory sizes: 1 – 16 million words (Word = 64 bits; multiply by 8 to get bytes)  
 (“Typical” system has 2 or 4 million words)
- Performance: Can sustain operation at hundreds of megaflops rate
- Require large mainframe serving as front-end

## Class VII: Next Generation

Scientific community is beginning to look at new areas of research never before possible. Very large memories are required. The more complex the problem, the greater CPU performance required to improve user productivity. Class VII supercomputer will become standalone system.

- Multiple CPU systems become necessary.  
 Each processor is a vector processing CPU (each is a supercomputer).
- Small number of CPUs  
 (Up to 8 in ETA-10)
- Memory size: up to 300 million words  
 (largest ETA-10 system has more than 300 million words.)
- Performance: Can sustain performance in gigaflop range  
 (Gigaflops – billions of floating points operation)  
 (ETA-10 has peak performance of 10 gigaflops.)
- Standalone system  
 not required to have front-end to access supercomputer– can assess it directly  
 (e.g., from workstation or a network)
- Supercomputers will see more interactive use instead of just batch processing.
- **ETA-10 is the first Class VII supercomputer.**

## In between VI & VII

- Multiprocessor Cray XMP – doesn’t have the memory of the next generation.
- Cray II – multiprocessor with large memory, but it’s debatable if its performance is in Class VII range.

**CLASS VI  
SUPERCOMPUTER**

---

**Single Processor**

**Vector Processing**

**1 – 16 MW Memory**

**100S Megaflops Performance**

**Front-end Required**

**CLASS VII  
SUPERCOMPUTER**

---

**Multi-processor**

**Vector Processing (in each CPU)**

**Up to 300 MW Memory**

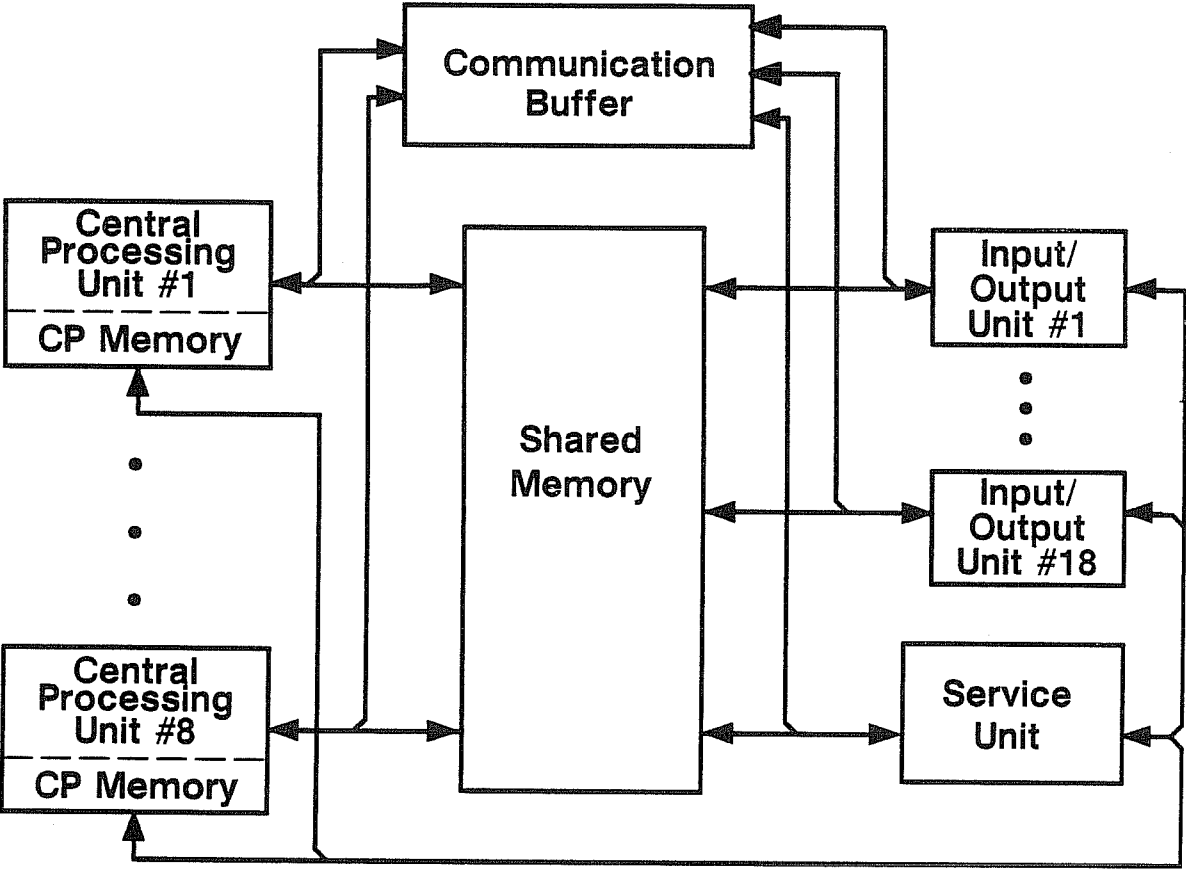
**Gigaflops Performance**

**Stand-alone Capability**

## ETA-10 FUNCTIONAL DIAGRAM

- **CPUs** – Up to 8 CPUs, each with its own memory of 4 million words.
- **CPUs** communicate with large shared memory – up to 256 million words.
- **Shared Memory** communicates with up to 18 Input/Output Units.
- **I/OUs** have channels going out to peripherals.
- **Service Unit** has operator console and maintenance console.
- **Communication Buffer** synchronizes events between CPUs and between CPUs and IOUs.
  - 1 million words of memory just for communication job
  - (Comparison: Some current supercomputers have 1 million words of total memory.)

# ETA<sup>10</sup> FUNCTIONAL DIAGRAM



## PICTURE OF MEMORY STACK

- 8 Cards/Stack
- Can be Shared Memory or CPU Memory
  - If Shared Memory:  
1 million words - 8 megabytes per stack
  - If CPU Memory:  
256K words per stack
- ETA-10
  - CPU: cooled in liquid nitrogen
  - Memory: air cooled

# PICTURE OF MEMORY STACK



# STORAGE DEVICE ROADMAP

Current production quantity/quality chips used in ETA-10:

- 64K Static MOS
  - In Central Processor Memory and Communication Buffer Memory
  - Speed needed here
- 256K Dynamic MOS
  - In large Share Memory

## GaAs

- Considered for logic devices
- Doesn't have high enough bit density to be considered for supercomputer memory

## ECL BiPolar

- Fast, but lower density
- Takes a lot of power
- Being replaced by MOS


## MOS

- Static MOS:
  - 64K bit on chip
  - Faster than dynamic
  - More expensive than dynamic
- Dynamic MOS:
  - 256K bits on chip

Note:

- Memory sizes are dramatically expanding.
- 1988 - Chips will be available with 4 times density.
- 256K - Static; 1 million bits - Dynamic
- ETA-10 is designed to take advantage of these new chips.
- ETA-10's chips will be able to be replaced with new chips that will quadruple the memory size.
- Then each CPU will have 16M 64-bit words of memory.
- Shared Memory can get as large as 1 billion words of memory

# STORAGE DEVICE ROADMAP



## LOGIC TECHNOLOGY ROADMAP 1985 – 87

This chart plots the switching speed vs. gate density. It shows why CMOS logic was chosen for ETA-10.

### Choices:

- Silicon ECL – current choice in all supercomputers
- GaAs – new technology – very fast
- CMOS (complementary metal oxide semiconductor)

**Switching speed** – Length of time it takes for logic gate to change its state

- GaAs – fastest
- Silicon ECL – slightly slower
- CMOS (at room temperature) quite a bit slower
- (Si CMOS semicustom currently available with 500 – 700 ps delay/gate)

However – if CMOS is cooled, the switching speed gets faster.

- If CMOS is cooled to liquid nitrogen temperature (80 degrees K), switching speed is doubled (to 300 – 400 ps)
- Still slightly slower than GaAs and silicon ECL but fast enough for a supercomputer.

Remember that switching speed is not the only factor determining speed of computer.

**Gate Density** – determines number of chips needed to make CPU for supercomputers

- |                                      |                                     |
|--------------------------------------|-------------------------------------|
| ● GaAs – 200 (very low chip density) | 10,000 chips needed                 |
| ● ECL – few thousand                 | few thousand chips needed           |
| ● CMOS – 20,000                      | fewer than 250 chips needed per CPU |

Density dramatically reduces chip connectors per CPU (200K – 100K connections).

ETA-10 uses only 240 chips.

Many fewer interconnections are needed using CMOS instead of GaAs.


### Reliability

- Number of interconnections directly related to reliability.
  - The fewer the interconnections, the greater the reliability.
- When CMOS is cooled, reliability is improved even more.

Power requirement for CMOS is very low and gives off very little heat.

**CMOS is the choice for ETA-10. All chips can be put on single PCB on ETA-10**

# LOGIC TECHNOLOGY ROADMAP 1985-87



# LOGIC TECHNOLOGY ROADMAP 1985 – 87, 1990 – 92

This chart shows today vs. next generation in logic technology.

## GaAs HEMT

- Switching time 20 – 50 ps/gate at liquid nitrogen temp.
- 6K – 20K gates per chip

Note:

GaAs HEMT logic used in the next generation of GaAs devices is very different from the ECL logic used in the current GaAs devices. A complete redesign of the chips will be required to go from GaAs ECL to GaAs HEMT so ETA has not gotten behind by not using GaAs today.


Note:

GaAs is still largely experimental. Its use in supercomputers has not been demonstrated yet.

## Si CMOS semicustom

- Higher switching speed but higher density, thus requiring fewer chips and less interconnects


# LOGIC TECHNOLOGY ROADMAP 1985-1987 & 1990-1992



## **LOGIC TECHNOLOGY ROADMAP 1990 – 92**

- Currently, it is unclear which will be the better route: higher density or faster switching speed.
- ETA is free to choose the best path.

# LOGIC TECHNOLOGY ROADMAP 1990-92





# SEMICONDUCTOR LOGIC TECHNOLOGY

## Advanced Large Scale Integration (ALSI) 20K chip

2000 – 3000 gates/chip used for B.E.S.T.

- When operating at liquid nitrogen temperature, performance doubles, but capability to run in air means initial checkout can be done at room temperature.
- Diagnosis of problem can also be done in air.

## More than 230 I/O ports/chip

- Means you can bring in two 64-bit operands and take out one 64-bit result on each chip.
- Example:  
You can put an entire arithmetic operation on single chip.

## B.E.S.T.

- Difficult to use oscilloscope probe on chip with 20K gates.
- Custom logic for self-testing
- B.E.S.T. is a trademark of ETA Systems.

# SEMICONDUCTOR LOGIC TECHNOLOGY

## CMOS GATE ARRAY [ALSI – 20K]

- > 15,000 Usable Gates / Chip
- > 230 I/O Ports [64-Bit Parallel Operation]
- Built-in Evaluation and Self-Test [B.E.S.T™] and Clock
- Operates at Room Temp. and LN<sub>2</sub> Temp. [ ≈ 77 degrees K]

# ALSI – 20K

## BUILT-IN EVALUATION AND SELF TEST

### Chip-to-Chip Continuity

- Checks for proper operation of the interconnect paths between arrays (chips) using either a predetermined pattern or a random pattern.
- Wiring is typically what fails, not logic, so it is very important to be able to diagnose these problems.
- Error display shows:
  - Physical layout of the array tested, highlighting the failing arrays
  - Additional error information such as failed array's name, type, revision, location, and connecting pin number of array connected to each pin.

### Validate Chip Type

- Ensures that the right chip, including revision number is at the right place.

### Remote Single Chip Validation

- Test individual chips when in CPU.
- Single chip can be validated.
- All chips in the system can be validated concurrently.

### Isolation of Failing Replaceable Device

- B.E.S.T. identifies the failing device so the unit can be replaced and repairs done in St. Paul.
- Shortening the time to check out the boards and improving the ability to diagnose problems help keep costs down.
- Note: If chip fails, you must replace the CPU.

**ALSI – 20K**  
**BUILT-IN EVALUATION AND SELF TEST**  
**– SYSTEM FEATURES –**

- **Chip-to-Chip Continuity Test**
- **Validate Chip Type – Revision**
- **Remote Single Chip Validation**
- **Isolation of Failing Replaceable Device**

## PICTURE OF CHIP

- Chip approximately 1 cm square.
- Large square is B.E.S.T. on chip.
- Small white squares – I/O pads.
- Gate array
  - Chip vendor can manufacture chips in quantity except for the last 3 layers.
  - Then, based on ETA-10 design, the last 3 layers can be added, connecting gates underneath to produce individual chip types.
  - ETA-10 has 90 different chip types.

## PICTURE OF CHIP

## SUPERCOMPUTER PACKAGING IMPROVEMENTS

This slide is a summary of the benefits of using new technology for logic and memory.

- CMOS uses less power (100:1)
- Fewer connections (26:1)
- Number of gates is up (1:80)
- Memory capacity is up (1:13)
- Number of gates in ETA-10 CPU  $\approx$  gates on CYBER 205 CPU, approximately 3 million

# SUPERCOMPUTER PACKAGING IMPROVEMENTS

<u>CYBER 205 - 1980</u>	<u>ETA<sup>10</sup> - 1986</u>	<u>RATIO</u>
20,000 Watts	200 Watts	100:1
786,000	30,000	26:1
3,200	240,000	1:80
5,000	64K	1:13



# PACKAGING BENEFITS

## Reliability

- Fewer interconnects mean fewer of the difficult to solve interconnect problems.
- Requires less power
- Liquid nitrogen lowers temperature:
  - Lower operating temperatures mean components last longer
  - Less heat to cause friction.

## Cost

- Lower both component cost and production cost as a result of fewer components because of increased chip density and fewer wires to connect components.

## Performance

- Switching speed is not the only factor.
- High packaging density provides faster speed:
  - More logic on a single chip
  - The distance signals have to travel is reduced; therefore, faster speed.
- Lower temperature improves CMOS switching speed – gives twice the clock speed of air-cooled system.

## **PACKAGING BENEFITS**

- **Improved RELIABILITY**
  - Lower Interconnect Count
  - Lower Operating Temperatures
- **Improved System COST**
  - Low Component Count
  - Low Interconnect Count
- **Improved System PERFORMANCE**
  - High Packaging Density
  - Cryogenic Operation

## TRANSMISSION TIME VS INTERCONNECT MATERIAL

Besides memory and logic chips, the 3rd component of technology in the ETA-10 is the printed circuit board.

This chart shows how fast the signal can be moved from one point to another.

- The reference is the speed of light.
- The next fastest medium is a coaxial cable (used to connect two PCBs).
- Printed circuit board is slower.
- Silicon (used in chip) is slowest.

Because coax is 80% of the speed of light, you might think coax is the material of choice; however, speed is not the only factor. You must also look at how far you must move the signal. When you look at the distance required, you realize that distance is more critical than speed.

- Coax cable distance may be about 3 feet.
- PCB distance may be 3 inches.
- But on a chip, the signal must move only 2 mm.
- Time delay is the speed times the distance traveled.

Design objective: To build the densest chip possible and a highly populated board to keep distances to a minimum.

The fastest way would be to have the entire computer on a single chip. This is possible with microcomputers but not with supercomputers.

The next best thing is to have the entire supercomputer on a single PCB.

- This is what we have done with ETA-10.
- We can do this because of high density CMOS chips.
- We only need 240 chips to make supercomputer.

**TRANSMISSION TIME  
VS  
INTERCONNECT MATERIAL**

<u>MATERIAL</u>	<u>FACTOR PS */MICRON</u>	<u>DISTANCE (MICRON)</u>	<u>DELAY PS *</u>	<u>RATIO</u>
Silicon	$7.5 \times 10^{-2}$	2K	150	1
PC Board	$10^{-2}$	150K	1500	10
Coax	$5 \times 10^{-3}$	900K	4500	30
Light (Ref)	$3 \times 10^{-3}$			

\* PS = Picoseconds

## **PRINTED CIRCUIT BOARD REQUIREMENTS**

Interconnect nodes provide device to device connections and device to I/O port connections.

Interconnects are very minimal.

Because they are all on the PC board, if a problem occurs, the whole CPU is replaced.

ETA-10 Printed Circuit Board

- 42 layers, 1/4" thick
- 16" x 22"
- Only ETA manufactures this special type of PCB.

## **PRINTED CIRCUIT BOARD REQUIREMENTS**

- **More Than 80,000 Interconnect Nodes**
- **Total Interconnect Length –  
2KM of “Wires”**
- **Size  $\approx 0.25 \text{ m}^2$**

## **PICTURE OF ETA-10 PRINTED CIRCUIT BOARD**

- Gold portion represents 240 chips for ETA-10.
- Red portion represents the PCB of CYBER 205.
- All logic in the 205 is now contained on 2 chips in the ETA-10.
- It took 65 of these red PCBs to make the CYBER 205 CPU.
- It takes 1 of these gold PCBs to make the ETA-10.

## PICTURE OF ETA-10 PRINTED CIRCUIT BOARD



## PICTURE OF ETA-10 AND CYBER CPUs

This slide shows the effect of the reduced number of PCBs, etc. in the ETA-10.

- In back is one of several cabinets of the CYBER 205 CPU.
  - Note: All of the coaxial cables are connecting the PCBs.
  - The equivalent of this wiring is contained in the layers of the single PCB of the ETA-10 CPU.
- In front is the CPU of the ETA-10.
  - The ETA-10 CPU is immersed in liquid nitrogen.
  - Memory is air-cooled.  
(16 stacks with 1/4 million words each; total of 4 million words)

## PICTURE OF ETA-10 AND CYBER CPUs

## **PICTURE OF ETA-10, LLOYD THORNDYKE, CYBER 205**

Note:

If you are giving only a brief overview of technology, you can use just 3 slides:

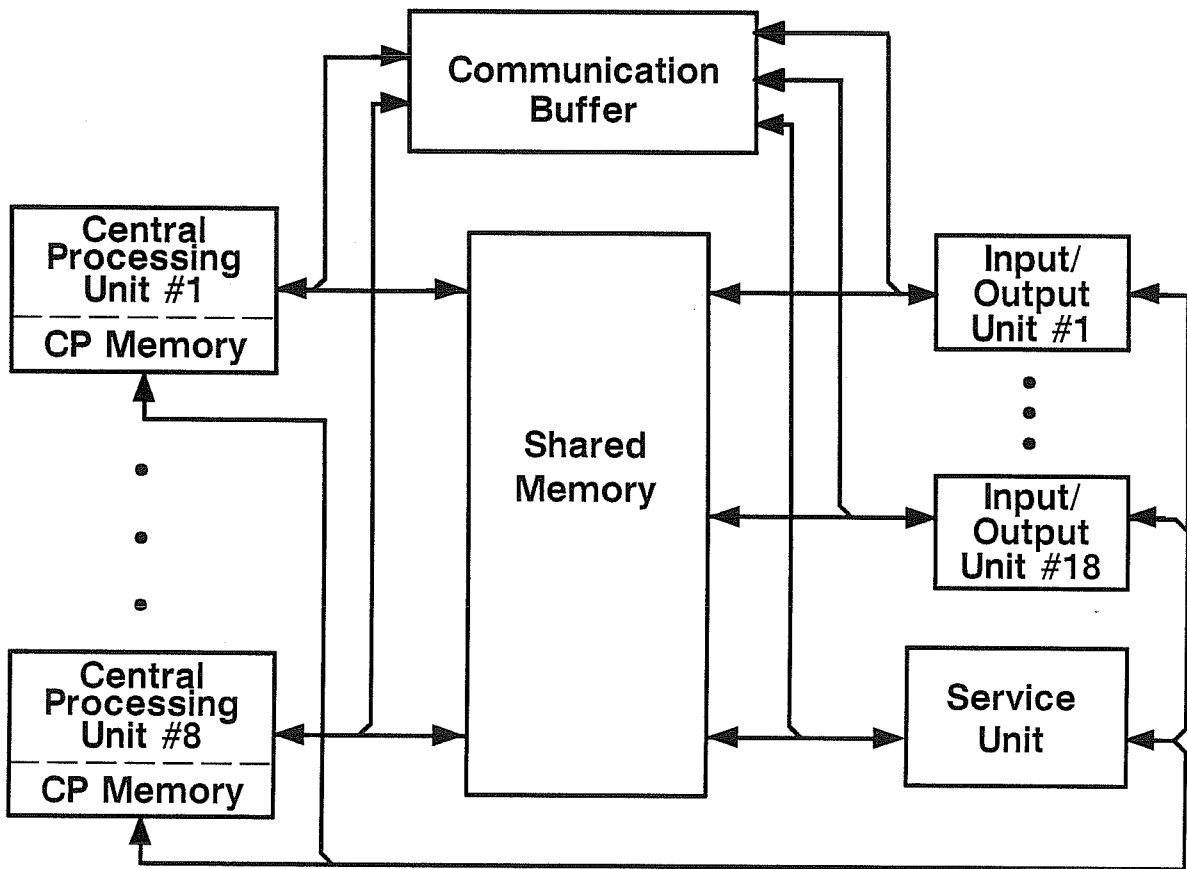
- Memory stack
- Chip
- Picture of ETA-10, Thorndyke, CYBER 205

# PICTURE OF ETA-10, LLOYD THORNDYKE, CYBER 205

## ETA-10 FUNCTIONAL DIAGRAM

- Up to 8 CPUs in ETA-10
  - Each CPU has 4 million words of CP memory.
- The CPUs communicate with a large Shared Memory (SM) which can be up to 256 million words.
- The SM also communicates with up to 18 Input/Output Units (IOUs) and to the Service Unit (SU).
- The IOUs have functional units that provide channels to peripheral devices and networks.
- The SU has operator and maintenance consoles.
- The Communication Buffer provides interprocessor communication.
  - It has 1 million words of memory.

# ETA<sup>10</sup> FUNCTIONAL DIAGRAM



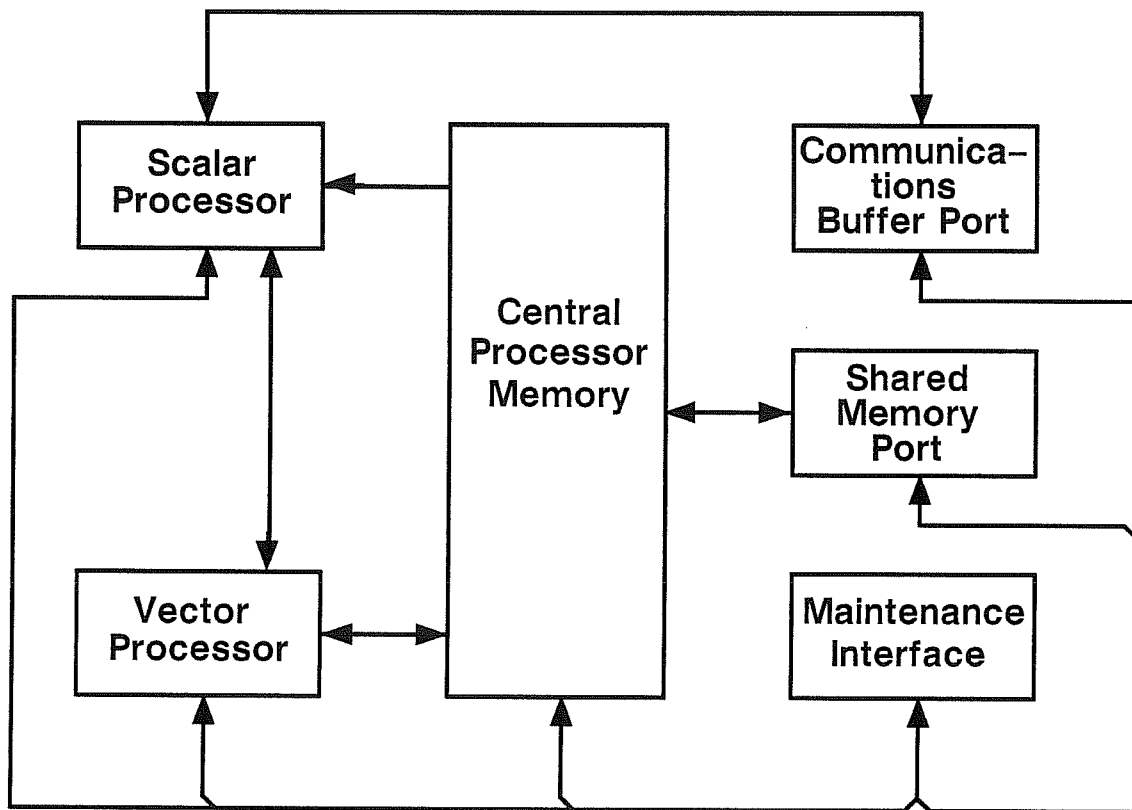
## ETA-10 CENTRAL PROCESSING UNIT

- CPU has:
  - A scalar processor and a vector processor that can operate simultaneously (i.e., they are independent of each other).
  - 4 million words of CP memory.
- Scalar and vector processors each have own path to CP memory providing for overlap of vector and scalar instructions.
- There are ports to the Communications Buffer (CB) and to Shared Memory (SM).
- Maintenance interface allows diagnostics to run in addition to doing performance monitoring.

### Note:

If you remove the SM and CB port, then the Central Processing Unit diagram is the same as for the CYBER 205 CPU.

# ETA<sup>10</sup> CENTRAL PROCESSING UNIT





# **ETA-10 SCALAR PROCESSOR CHARACTERISTICS**

## **Cyber 205 Compatibility**

- ETA-10 is instruction set compatible with CYBER 205.
  - i.e., All instructions on 205 are on ETA-10
  - There are additional instructions on the ETA-10 to handle additional CPUs, memory hierarchy, etc.

## **Instruction Issue Every Clock Cycle**

- ETA-10 has **scalar processor**
- ETA-10 issues new instruction every clock cycle.
  - This is important because some supercomputers take more than one cycle to issue an instruction.
  - It is important to know how much work is being done in each clock cycle and not just know clock cycle time.

## **64 Word Instruction Stack**

- Large instruction stack allows for large in-stack loops within the CPU.

## **236 Word Register File**

- Very large register file maximizes number of calculations done without having to store or fetch from memory.

## **Independent, Segmented Functional Units**

- These units allow multiple instructions to operate concurrently.

# **ETA<sup>10</sup> CENTRAL PROCESSOR SCALAR CHARACTERISTICS**

- **Cyber 205 Compatibility**
- **Instruction Issue Every Clock Cycle**
- **64 Word Instruction Stack**
- **256 Word Register File**
- **Independent, Segmented Functional Units**

# ETA-10 VECTOR PROCESSOR CHARACTERISTICS

## Vector Processor:

### Two Floating Point Pipelines

- ETA-10 is a memory to memory vector machine:
  - This means: You get 2 vectors in memory, send them to vector unit; resulting vector is streamed back to memory.
- With 2 floating point pipelines, once the vectors are started up, we get two 64-bit results every clock second

### 32-Bit Floating Point Arithmetic

- In 32-bit (half-precision mode), we get 4 results every clock second.
- IMPORTANT: We double the vector performance of the computer if we can do 32-bit arithmetic. If your algorithm or application can get away with 32-bit precision, then you can double your vector performance as well as double your memory size.
- There are some applications (e.g., weather models) that can utilize 32-bit arithmetic.

### Linkage Potential

- Hardware exists to link a scalar operation to a vector operation and perform it as one operation. This doubles operations performed per clock cycle.
- Example: You can multiply a vector times a scalar and then add it to another vector. You can do that in the time it would take to do one operation. (You've done two operations for the price of one.)

### String Unit

- This allows vectors to be processed as individual string of bits rather than as a 32-bit or 64-bit word.
- Bit strings are used primarily as control vectors. They control data motion in memory so we can make more effective use of the vector unit.

### Vector Shortstop Path

- New on ETA-10
- If the resulting vector that goes back to memory is used in a subsequent vector operation, you may not want to wait for it to go back to memory.
- Hardware on ETA-10 lets you send that vector right back into the vector pipeline.
- This improves the vector startup time and gives better short vector performance.

## **ETA<sup>10</sup> CENTRAL PROCESSOR VECTOR CHARACTERISTICS**

- **Two Floating Point Pipelines**
- **64-bit and 32-bit Floating Point Arithmetic**
- **Linkage Potential (64-bit, 32-bit)**
- **Overlap of CPU Operations with I/O**
- **Overlap of Scalar with Vector**
- **String Unit**
- **Vector Shortstop Path**

# CPU MEMORY CHARACTERISTICS

## 64K MOS SRAM

- It is extremely fast.

## Four Million 64-bit Words

- **Expandable with technology** means when the new 256K static RAM chips are available at a production level, you can replace these chips and have 16M words of memory. (You can swap memory boards and expand memory by a factor of four.)

## SECED

- (Single Error Correction - Double Error Detection)
- Error correction on each 32-bit half-word lets us correct significantly more memory errors than if we just had error correction on 64-bit full-words.

## 75 Billion Bits/Sec Bandwidth

- Between CP memory and CPU
- Can provide data to CPU as fast as CPU can process it
- High enough to sustain both vector and scalar units at their maximum performance.

## Virtual Addressing

- Easy porting of codes to and from other machines
- Easy for programmers
- Memory structure transparent to the user
- Unlimited address space.

# **ETA<sup>10</sup> CPU**

## **MEMORY CHARACTERISTICS**

- **64K MOS SRAM**
- **Four Million 64-bit Words  
(Expandable with Technology)**
- **SECDED for Each 32 Bits**
- **75 Billion Bits/Sec Bandwidth**
- **Virtually Addressed**

# COMMUNICATION BUFFER

Communication Buffer provides communication between all CPUs, IOUs, and SUs to coordinate multiprocessing activities, data fetching and storing, and system table updating.

## Queuing Operations

- A flag is set in the CB indicating there is work to be done.
- Example:
  - A job comes into ETA through IOU.
  - Control language and data are put into Shared Memory.
  - Additional information is put into the Communication Buffer.
  - If a CPU has resources available, it can quickly check with the CB to see if there is additional work to be done.
  - If CB indicates there is more work, it can get enough information to get the job and bring it into the CPU.

## Semaphore Operations

- Used when you need to coordinate some type of shared resource
- Signifies whether some type of event has occurred or not (e.g., locking a system shared table)
- Example:
  - Multitasking job where two tasks are running on separate CPUs but they're sharing some data in SM.
  - If you also assume that they're writing to that shared data, you must synchronize the events.
  - Communication Buffer provides that synchronization.
  - The CB gives you rapid access to small amounts of data.
  - The CB actually communicates with the register file in the CPUs.

Communication Board has same memory as CPU (expandable with technology).

Processor in CB is Communication Buffer Interface.

## **ETA<sup>10</sup> COMMUNICATION BUFFER**

- **Low Overhead Interprocessor Communications**
- **Queuing Operations**
- **Semaphore Operations**
- **One Million 64-Bit Words CB Memory (Expandable with Technology)**
- **Communication Buffer Interface Board**



## COMMUNICATION BUFFER INTERFACE

- Manages all of the communication and coordination
  - Uses 20,000 gates CMOS chips on 42-layer PCB
  - Board is air-cooled - gives sufficient speed ,
- Uses same chip technology, high-tech PC board as the CPU
- Almost as much logic in both the Communication Buffer Interface and the Shared Memory Interface as in the Cray 1 supercomputer.

# **ETA<sup>10</sup>**

## **COMMUNICATION BUFFER INTERFACE**

- **Uses ETA<sup>10</sup> CMOS Chip Technology**
- **42 Layer PC Board**
- **Air Cooled**
- **System-wide Control between CPUs, I/O Processors, and Service Unit**

## SHARED MEMORY CHARACTERISTICS

- Slower dynamic RAM but much denser
- As technology advances and the 1 million DRAM chip becomes available in reliable quality and quantity, we can simply replace memory boards. The system will then have 4 times the SM capacity or up to 1 billion words.

### Eight Hi-speed CPU Ports

- One of these ports for every CPU on the system
- Very high speed (10 billion bits/sec.)
- The fact that each CPU has its own Shared Memory port distinguishes Shared Memory from a solid state disk. On a solid state disk, you might have one or two ports that must be shared among all of the CPUs in the system.
- In the ETA-10, all CPUs have their own ports to Shared Memory and all can be driven simultaneously at maximum rates  
(10 billion bits/sec. per port – 80 billion bits/sec. total)

### Eighteen Lo-speed Ports

- To I/O units
- Each low speed port can simultaneously transfer data between SM and I/O units at 440 million bits/sec. per port – 7920 million bits/sec. total.

Processor in SM is Shared Memory Interface (SMI) board.

## **ETA<sup>10</sup>**

### **SHARED MEMORY CHARACTERISTICS**

- **256K MOS DRAM**
- **32 – 256 Million Words with SECDED  
(Expandable with Technology)**
- **Eight Hi-speed CPU Ports  
(10 Billion Bits/Sec Bandwidth per Port)**
- **Eighteen Lo-speed I/O Ports  
(440 Million Bits/Sec Bandwidth per Port)**
- **Shared Memory Interface Board**

## SHARED MEMORY INTERFACE

- Same technology as CPU and CBI boards
- Uses 20,000 gate CMOS chips on 42-layer PCB
- Air cooled
- Initiates transfers between SM and CPM or IOU
- Maintains Shared Memory tables
- Provides lockout features
- Coordinates sharing of code and data within SM
- Provides control to move data between CPU and Shared Memory and between Shared Memory and I/Os
- Nearly as much logic in the Shared Memory Interface as there was in the Cray I
- Because of Shared Memory Interface, the CPUs themselves don't have to worry about this data motion.

## **ETA<sup>10</sup> SHARED MEMORY INTERFACE**

- **Uses ETA<sup>10</sup> CMOS Chip Technology**
- **42 Layer PC Board**
- **Air Cooled**
- **Provides Access to Shared Memory**

## **ETA INPUT/OUTPUT UNIT**

### **Shared Memory**

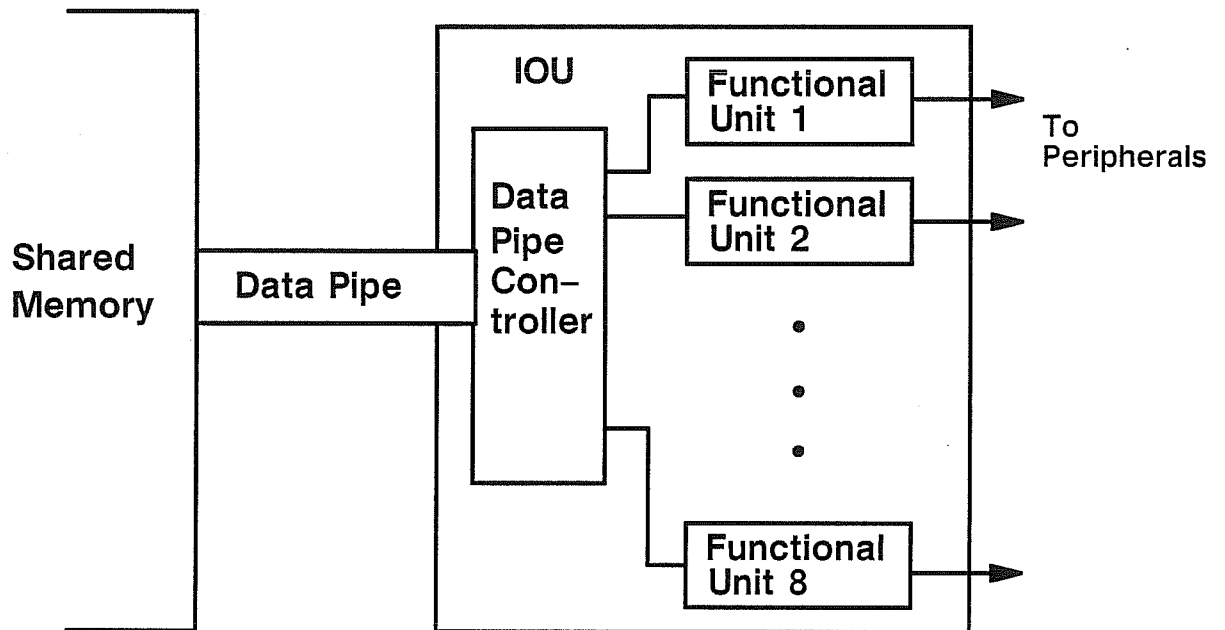
### **Data Pipe**

- Consists of four 110 megabit fiber optic cables with a peak data transfer rate (each direction) of 440 megabits/sec at 100 feet.
- Fiber Optic Link can be up to 1000 feet long to permit you to move the IOUs and peripherals away from the Shared Memory.

### **IOU**

- Data Pipe Controller
- Up to 8 functional units depending on the type of peripherals attached
- Each functional unit is a computer – a microprocessor with its own memory.

# ETA<sup>10</sup> INPUT/OUTPUT UNIT (IOU)





## **IOU FUNCTIONAL UNITS – TYPICAL CONFIGURATIONS**

### **Data Management Processor (DMP)**

- Single board computer with processor, memory and Data Pipe buffer
- Manages disk I/O.

### **Disk Channel Controller (DCC)**

- Single board computer
- Controls access to disk
- If you have this in IOU, you need disk management processor.
- Provides 96 megabits/sec channel out to disks.

### **FIPS (FCC)**

- Single board computer
- Provides FIPS Standard Interface for tapes
- LCN also uses FIPS Standard Interface.
- IBM compatible channel
- 24 megabit/sec channel goes either to tapes or networks.

### **VME (VCC)**

- Single board computer
- Conforms to VME bus interface standard
- Used on ETA-10 to provide access to Ethernet network
- Rated at 10 megabits/sec.

Can configure up to 2 DCCs in single IOU

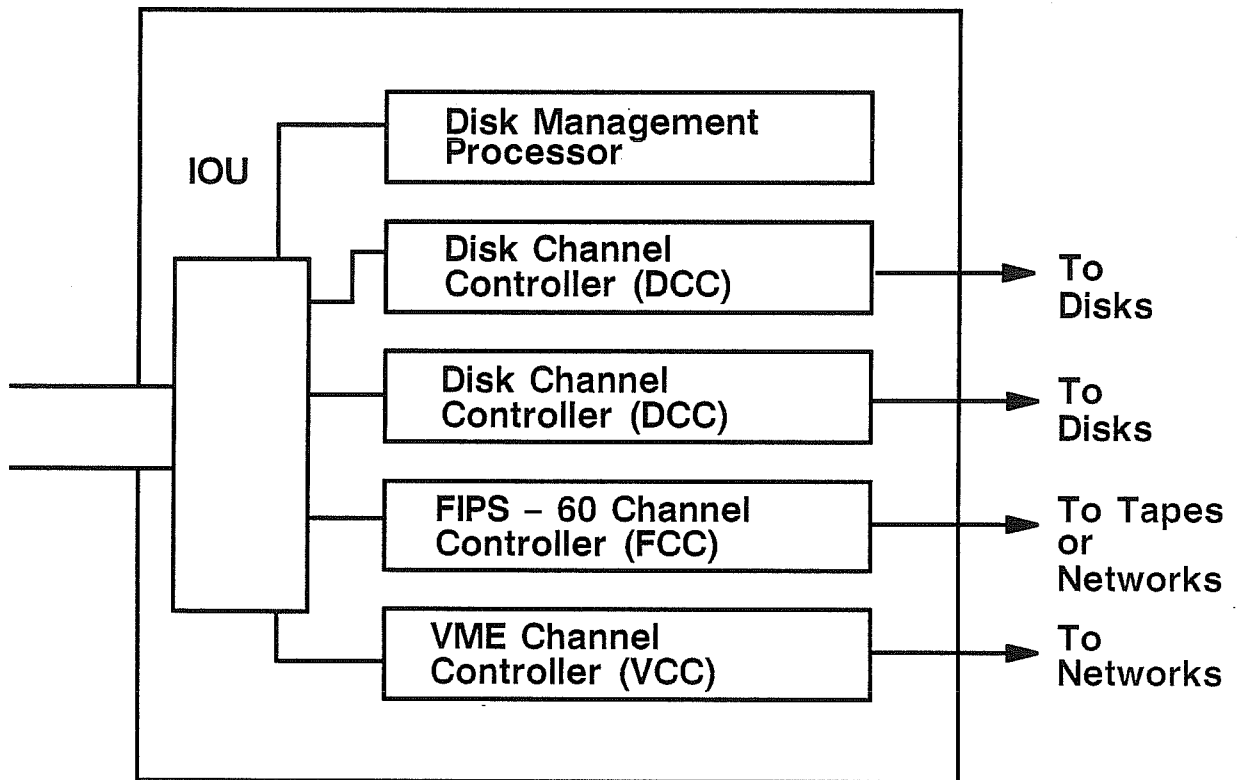
2 DCCs allow other 6 slots to be filled with FCCs and/or VCCs  
(DMP required with DCC).

If there are no DCCs in the IOU, you can use all 8 slots for FCCs and/or VCCs  
(no DMP is required).

# ETA<sup>10</sup>

## IOU FUNCTIONAL UNITS

### TYPICAL CONFIGURATIONS



## ETA-10 DISKS

- Hydra disks from Control Data or Ibis disks can transfer from 4 disks on 1 IOU concurrently
- Disk transfer rate:
  - 96 megabits/sec. peak (12 megabytes)
  - 80 megabits/sec. sustained
- Highest performing disks on market today
  - 12 megabyte/sec transfer rate (96 bits/sec)
  - IO unit can go up to 440 megabits/sec
  - We can have more than one disk channel going through an I/O unit.
- Up to 60 simultaneous transfers from disks to the ETA-10
  - Very high I/O bandwidth, such as we have on the ETA-10, is extremely important for supercomputers.
- We can have large number of disks configured on ETA-10 to get a very high total mass storage.
- Assuming 18 IOUs with 2 DCCs per IOU, 8 disks per DCC
  - Mass storage capacity is 345 gigabytes.

## **ETA<sup>10</sup> DISKS**

- **1200 Megabyte/Unit Capacity**
- **12 Megabyte/Second Transfer Rate**
- **Up to 60 Simultaneous Transfers**
- **On-line Mass Storage to 345 Gigabytes**

## ETA-10 TAPES

- Standard high performance tapes
- Up to 200 tape drive connections
  - We offer this to meet the requirements of certain users, like those doing seismic processing, who have many tapes and can't afford the delay in performance by staging it through a front-end.
- Most systems will have only a few tapes for file archiving.

## **ETA<sup>10</sup> TAPES**

- **9-Track, 200 Inches per Second**
- **6250 BPI GCR**  
**1250K Bytes/Second**
- **1600 BPI Phase Encoded,**  
**320K Bytes/Second**
- **Up to 200 Tape Drive Connections**
- **ANSI Standard Labels**

## **ETA-10 NETWORKS**

### **High Speed Network**

- Control Data Corp. LCN and NSC HYPERchannel
- 50 megabit/sec rated network primarily for file transfer between mainframes.

### **Open Interconnection Network**

- 10 megabits/sec for interactive access to the ETA-10.

# ETA<sup>10</sup> NETWORKS

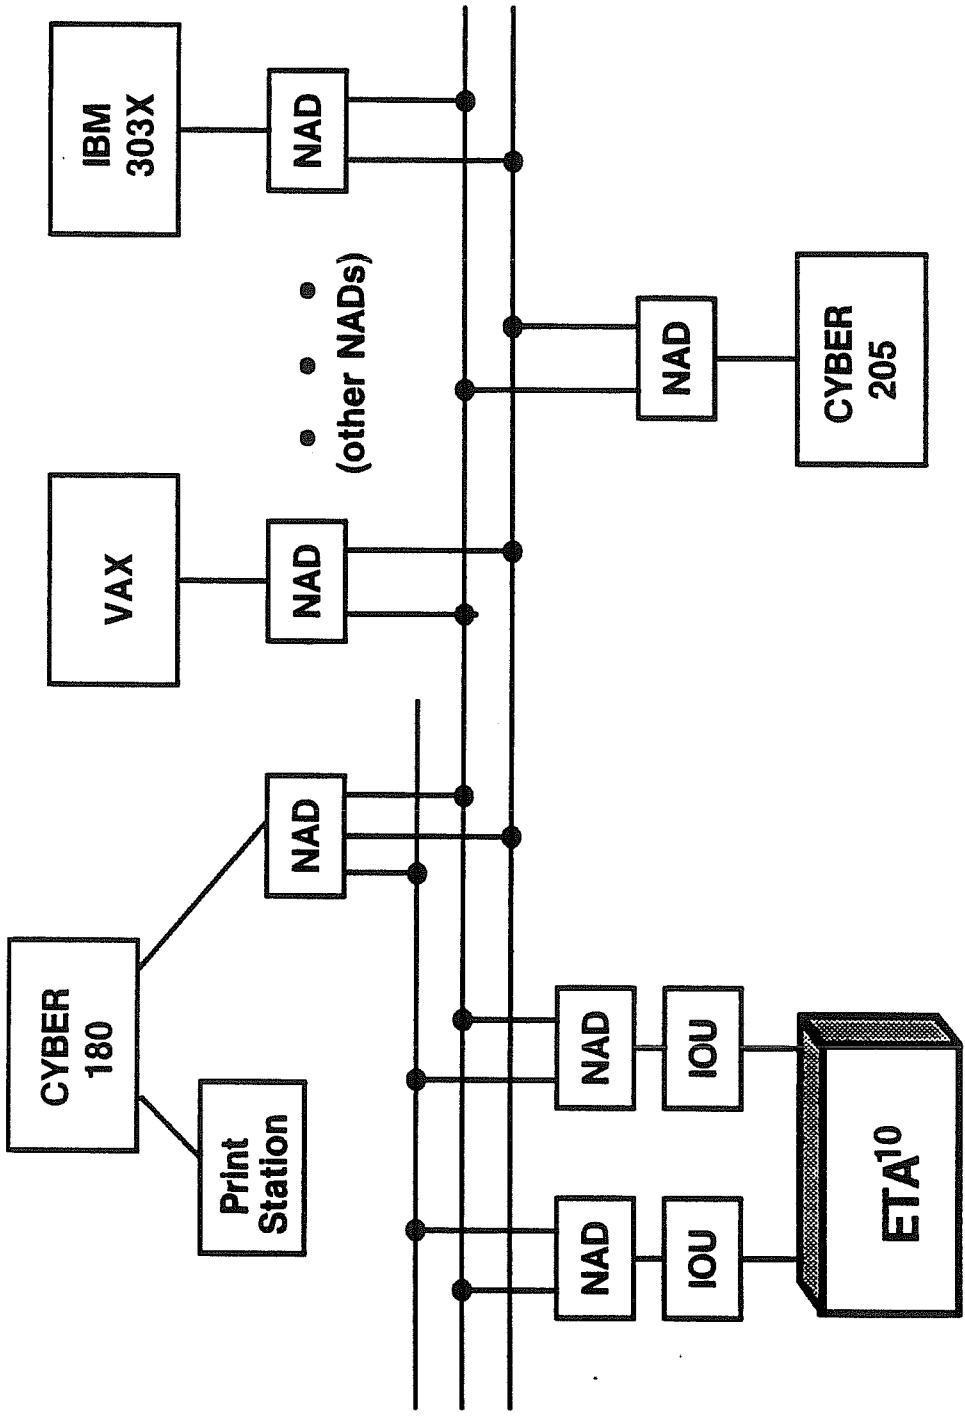
- **High Speed Network**
  - Control Data Loosely Coupled Network (LCN)
  - Hyperchannel Planned
- **Open Interconnection Network**
  - Ethernet Running TCP/IP
  - Open Systems Interconnection (OSI) Planned



## **LOOSELY COUPLED NETWORK**

- NAD – Network Access Device  
(The CYBER 205 is also connected to other computers in this way.)
- Primary purpose of network: very high speed file transfers
- Typically used for batch processing.

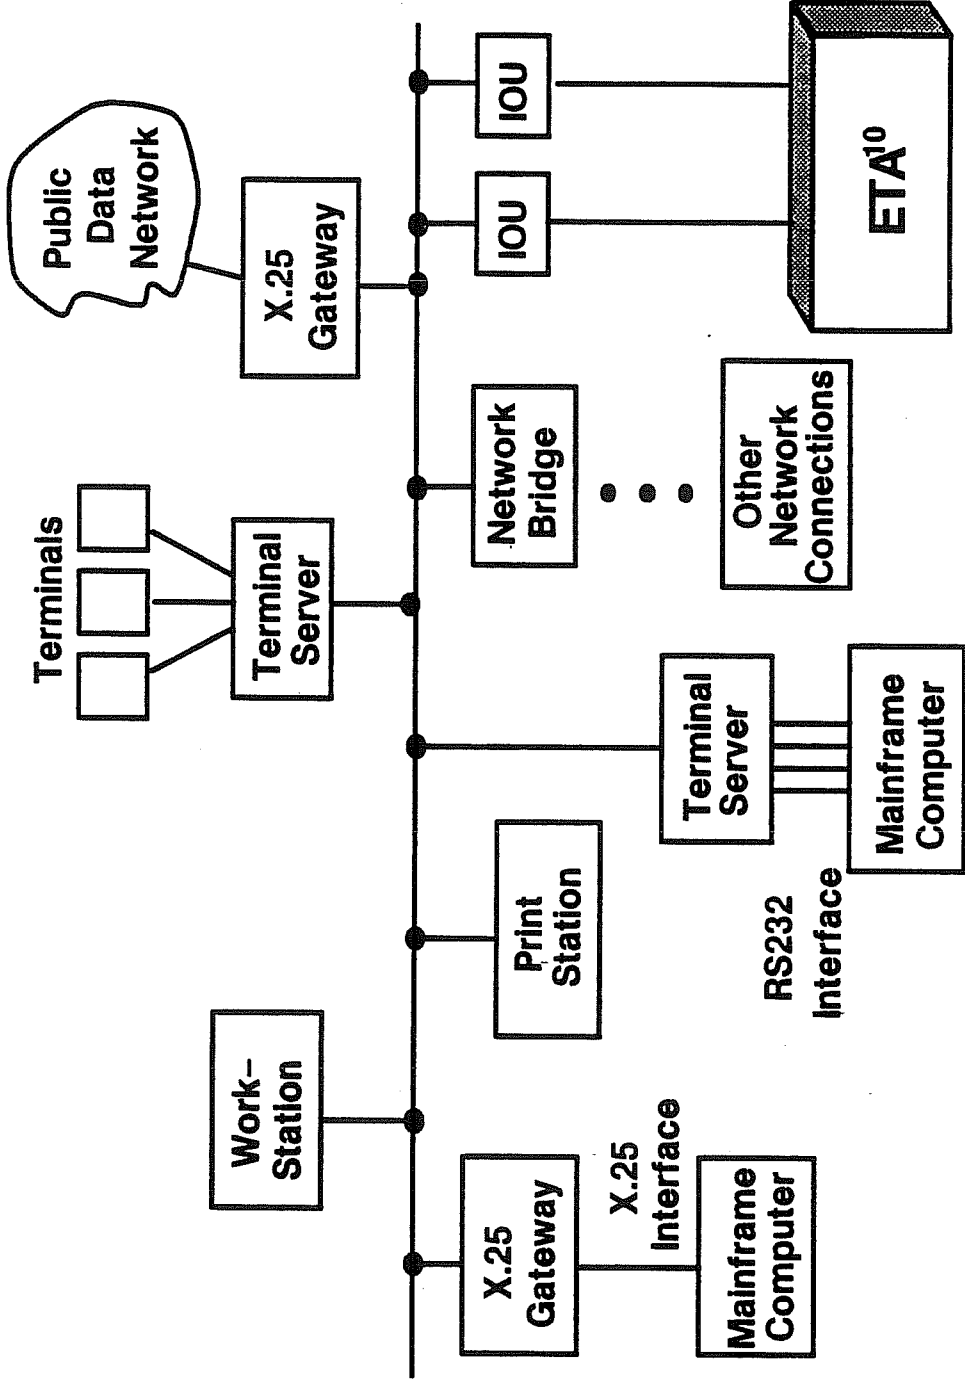
# LOOSELY COUPLED NETWORK



## OPEN INTERCONNECTION NETWORK

- ETA-10 is connected to Ethernet through the IOU.
- Workstation has direct access to the ETA-10.
- Terminals through terminal servers, gateways to other networks
- Only way that we will do interactive processing is through the open interconnection network to the ETA-10

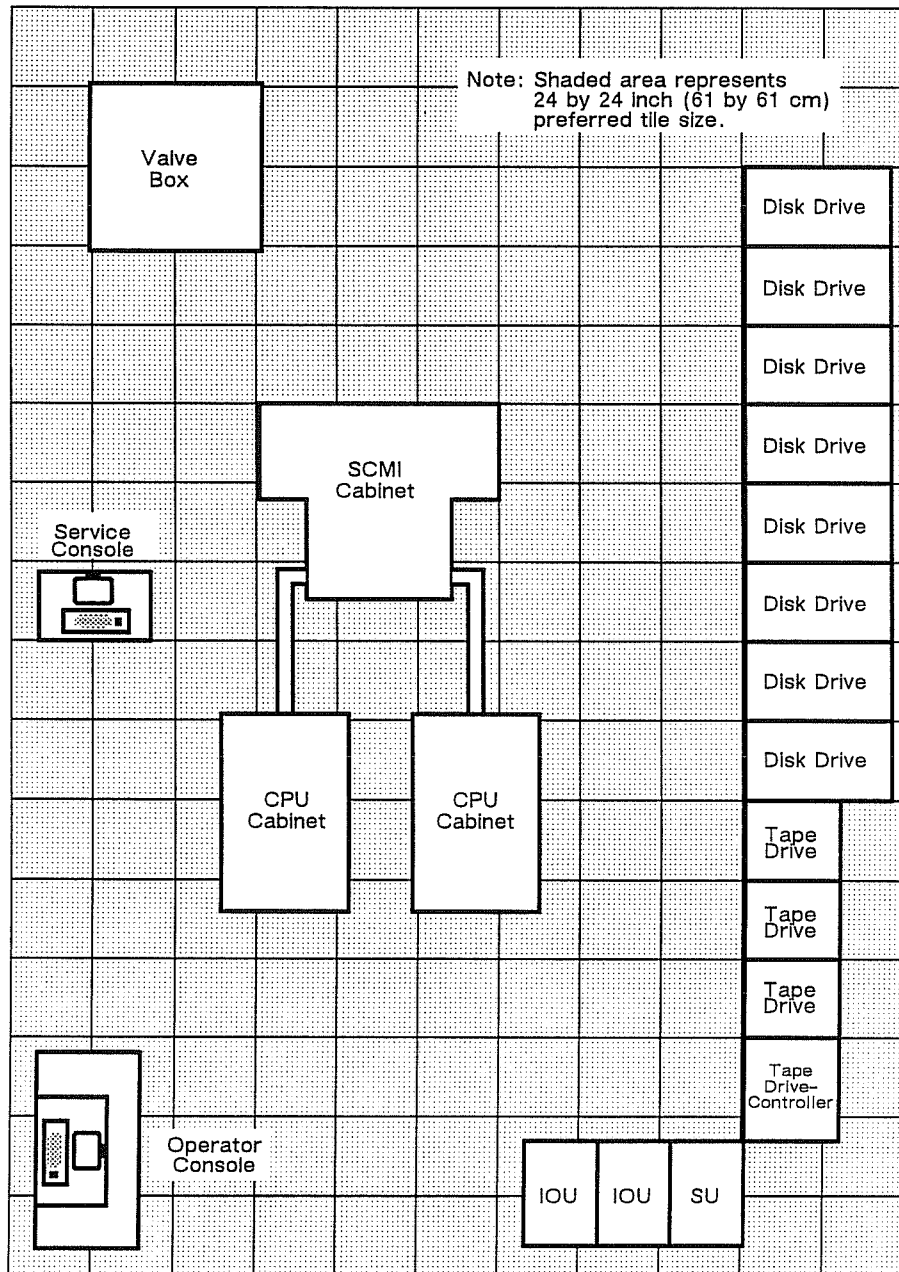
# OPEN INTERCONNECTION NETWORK



## **TYPICAL MAINFRAME AND PERIPHERAL EQUIPMENT**

- 2 - 4 CPU configuration
- **Shared Memory** - full 256K words
- **Service Console** - to run diagnostics from and connect to RSS
- **Operation Console** - operator and system administrator environment

# TYPICAL MAINFRAME AND PERIPHERAL EQUIPMENT LAYOUT



ETA SYSTEM

## ETA-10 AVAILABLE REDUNDANCY

- “Degradability” might be a better word.
- ETA-10 is designed to be most reliable supercomputer in world.
- Design objective: To make a supercomputer available to the user 100% of the time.
  - Therefore, there is the need to address possibility of hardware component failing.
  - We have at least 2 of every kind of component in ETA-10.
- Example: CPU in ETA-10:
  - Inserted into cryostat containing liquid nitrogen
  - Cryostat has two slots, each has CPU.
  - If one CPU fails and must be fixed, your whole system must go down because you must warm up the entire cryostat and take out the CPU.
  - To avoid this, ETA uses two cryostats, each with one CPU.
  - If one CPU fails, take it out, and replace it while the other good CPU continues to operate.
  - The system is never more than 50% unavailable to the user.
- Not everyone requires this redundancy –  
For a site that doesn’t need it, ETA can offer a lower priced system.

### Dual Shared Memories

- Shared Memory is divided into two halves.
  - Each half has its own SMI.
  - Each half has a high speed port to each of the 8 CPUs and 10 low speed ports (9 to IOUs and 1 to SU).
- Need two Shared Memory Interface boards
  - If one fails, Shared Memory can continue to operate with one good interface board.

### Dual Communication Buffers (Dual Comm. Buffer Interface Boards)

- Memory is divided into two half-million word partitions, each with its own CBI.
- SM and CB maintain duplicate tables in both halves of their respective memory.
- Depending on the critical nature of the table, the frequency and volume of updates required determines if and how a copy is maintained.

### Dual I/O Paths

- By configuring each disk with access to two IOUs, no single IOU failure would cause loss of access to any data.

### Dual Power Supplies

- Dual power supplies require so little power, UPS is feasible.

### Either Half Maintenance Philosophy

- Failed component can be configured out of the system with software, allowing maintenance to be done and the component returned to the system without the user’s knowledge

## **ETA<sup>10</sup> AVAILABLE REDUNDANCY**

- **Dual Shared Memories**
- **Dual Communication Buffers**
- **Pairs of CPUs**
- **Interlocked Software Control**



## **ETA<sup>10</sup> AVAILABLE REDUNDANCY**

- **Dual I/O Paths**
- **Dual Power Supplies**
- **Dual Cooling Systems**
- **Either Half Maintenance Philosophy**

## **ETA-10 CONFIGURATION OPTIONS – E SERIES**

- E Series is a lower cost, lower performance system than G series.
  - 4 processor E series will still be most powerful supercomputer around (except for G series)
- E Series is not offered with redundancy (degradability) option.
  - Only one Shared Memory and one Communication Buffer
- E Series will have clock cycle of 10.5 nanoseconds.
- E Series is field upgradable to the G series.

# **ETA<sup>10</sup> CONFIGURATION OPTIONS E SERIES**

- **Processor Options**
  - 1, 2, 3, 4 CPUs with 4 Million Words CPU Memory Each
- **Shared Memory Options**
  - 32, 64, 96, 128 Million Words
- **I/O Options**
  - 2 to 9 I/O Units
- **Performance**
  - Up to 3.5 Gigaflops

## **ETA-10 CONFIGURATION OPTIONS – G Series**

- G Series is a higher performance, higher cost series.
- G Series will have clock cycle of 7 nanoseconds.
- Redundancy (degradability) features are not necessary for the functionality of the 2 and 4 processors system; therefore, they are optional.
- For 6 and 8 processor system, dual SMI boards and dual CBI boards are required for functionality; therefore, they are standard.
- Performance – full 8 processor system can get up to 10 gigaflop peak rate.

# **ETA<sup>10</sup> CONFIGURATION OPTIONS G SERIES**

- **Processor Option**
  - 2, 4, 6, 8 CPUs with 4 Million Words CPU Memory Each
- **Shared Memory Options**
  - 64, 128, 192, 256 Million Words
- **I/O Options**
  - 2 to 18 I/O Units
- **Optional Redundancy on 2 and 4 Processors**
  - Multiple Cryostats
  - 2 SMIs, 2 CBIs
- **Redundancy Standard on 6 and 8 Processors**
- **Performance**
  - Up to 10 Gigaflops

## Picture of ETA-10 (4 Processor System)

- Front cabinet
  - Cryostat - each cryostat can have two CPUs.
  - CPU memory is air cooled.
  - Power supplies
    - Run on 50-60 cycle power
    - No Motor Generator (MG) set is required.
  - For a 6 or 8 processor, you need more of the front cabinets.
- Back cabinet
  - Shared Memory
  - Shared Memory Interface
  - Communication Buffer Memory
  - Communication Buffer Interface
  - Interface for fiber optics cables
- In back room (not shown)
  - Cryogenerator
    - (Cryogenerator and holding tank that contains a reservoir of liquid nitrogen).
  - Reservoir has enough liquid nitrogen to keep ETA-10 running for 6-8 hours. This time allows you to fix cryogenerator or to get more liquid nitrogen.
  - Chillers for air conditioning used to cool memory, SBI and CBI boards.

**THANK YOU**