



Overview

Products

EV6

Tsunami & Typhoon Chip Sets

Dennis Hayes for Goldrush Team

- **Products**
- **Basic Architecture**
- **EV6 - 21264 Features**
- **Functional Description and Block Diagrams**
 - **Chip Types**
 - **Interfaces: CPU, Memory, I/O**
- **Probes**
 - **CPU**
 - **I/O**
- **Questions and Answers**


Products

- **Monet - Uni EV6 Workstation**
- **Goldrush - Dual EV6**
- **Catamaran - Dual EV6**
- **Clipper - Quad EV6**
- **Wildfire - Turbolaser class machine**
- **Turbolaser EV6**

Roll Out

- ✓ Monet, Goldrush, Catamaran and Turbolaser-EV6 will be the first implementations of the EV6 architecture.
- ✓ Goldrush is initially targeted at Customers who require an EV6 platform for evaluation and application porting.
- ✓ Catamaran will be the first EV6, high density, 5U rack mount system in the Compaq ProLiant Series .

Block Architecture



- One to four CPUs
- Two 1.2GB/s switched CPU/memory connections
- 2.4GB/s aggregate memory bandwidth
- 4.8GB/s aggregate system BW
- Up to 8GB, 4GB per memory board
- 1-2 64b PCI buses (10 slots)
 - 512 MB/s I/O BW (full-duplex)
- No slot trade-off for CPUs, memory, or I/O

- **3 Planned Implementations:**

- EV6, 0.35-micron, six-layer-metal CMOS process 2.2-volt core.
- EV6/7, 0.25 micron and EV6/8 0.18 micron performance enhancements and cost reductions.

- **Architecture**

- Superscalar, deeply pipelined design
- Four instructions per clock cycle to be issued to four integer execution units and two floating point units.
- Enhanced instructions include specially developed Motion Video Instructions (MVI)

- **Features:**

- Out-of-order instruction execution
- Large (64KB) on-chip data and instruction caches
- Improved branch prediction through intuitive execution
- Increased bandwidth for high-speed access to level 2 cache and system memory.

Internal Design

- It contains four integer execution units, two of which can perform memory address calculations for load and store instructions. It also contains dedicated execution units for the floating-point add, multiply, divide and square root operations.
- The on chip instruction cache Icache is a 64KB, two way set-associative, virtually addressed cache with 64 byte blocks. The on chip data cache Dcache is a 64KB, two way set-associative, virtually indexed, physically tagged, write back cache with 64-byte blocks.

External Interface

- Two ports
 - Bcache port with a 16-byte data bus
 - System port with an 8-byte data bus.
- Bcache port is used, on Goldrush, to support a 4 MByte level 2 backup cache Bcache. The EV6 completely controls all Bcache operations.

Internal sections:

- Integer execution unit Ebox
- Floating point execution unit Fbox
- Instruction fetch, issue and retire unit Ibox
- Memory reference unit Mbox
- External cache and system interface unit Cbox
- Instruction cache Icache
- Data cache Dcache

Performance Estimates:

- 500-600 MHz Operating Frequency - EV6
- 40 SPECint95, 60 SPECfp95
- MPEGII Encode in real time
- 4+ GB/s level 2 Cache, 2+ GB/s Memory Peak
- 1600 MB.s McCalpin STREAM
- 900 MFLOPS Linpack 1000 x 1000

Tsunami and Typhoon

The Tsunami and Typhoon chip sets implement a switch based architecture connecting the CPU's, Memory and I/O sub systems. Tsunami supports up to 2 processors and Typhoon up to 4.

Three chip types are used to implement the switch:

- **C-Chip**, C2 for Tsunami and C4 for Typhoon, provides the interface from the CPU and main memory. It also includes an interface for flash memory and interrupts, the TigBus Interface. One C-Chip is used per system
- **D-Chip** provides the data path connecting the CPU, Memory and I/O. Systems may use two, four or eight D-Chips.
 - Clipper and Goldrush uses eight D Chips to provide two 256 bit memory interfaces.
- **P-Chip** provides an interface to a PCI Bus. One or two P-Chips may be used depending on the system design.
 - Clipper and Goldrush uses two P-Chips to support two 64 bit PCI Busses.

The Tsunami chipset has the following design attributes:

- up to two EV6 processors without using duplicate tags
- up to two 64 bit, 33mhz PCI buses, each with its own PCI address space
- support for Synchronous DRAM main memory, in 16Mb and 64Mb technology
- support for a large range of main memory capacity (16MB to 1GB using 16Mb DRAMS, up to 4GB using 64Mb DRAMS)
- support for ECC in main memory
- system clock periods from 10ns to 15ns
- low latency memory access (100ns CPU access using 100Mhz synchronous DRAMs)
- very high bandwidth (3.2GB/s peak memory bandwidth per processor using 100Mhz synchronous DRAMs)

Typhoon - Tsunami Differences

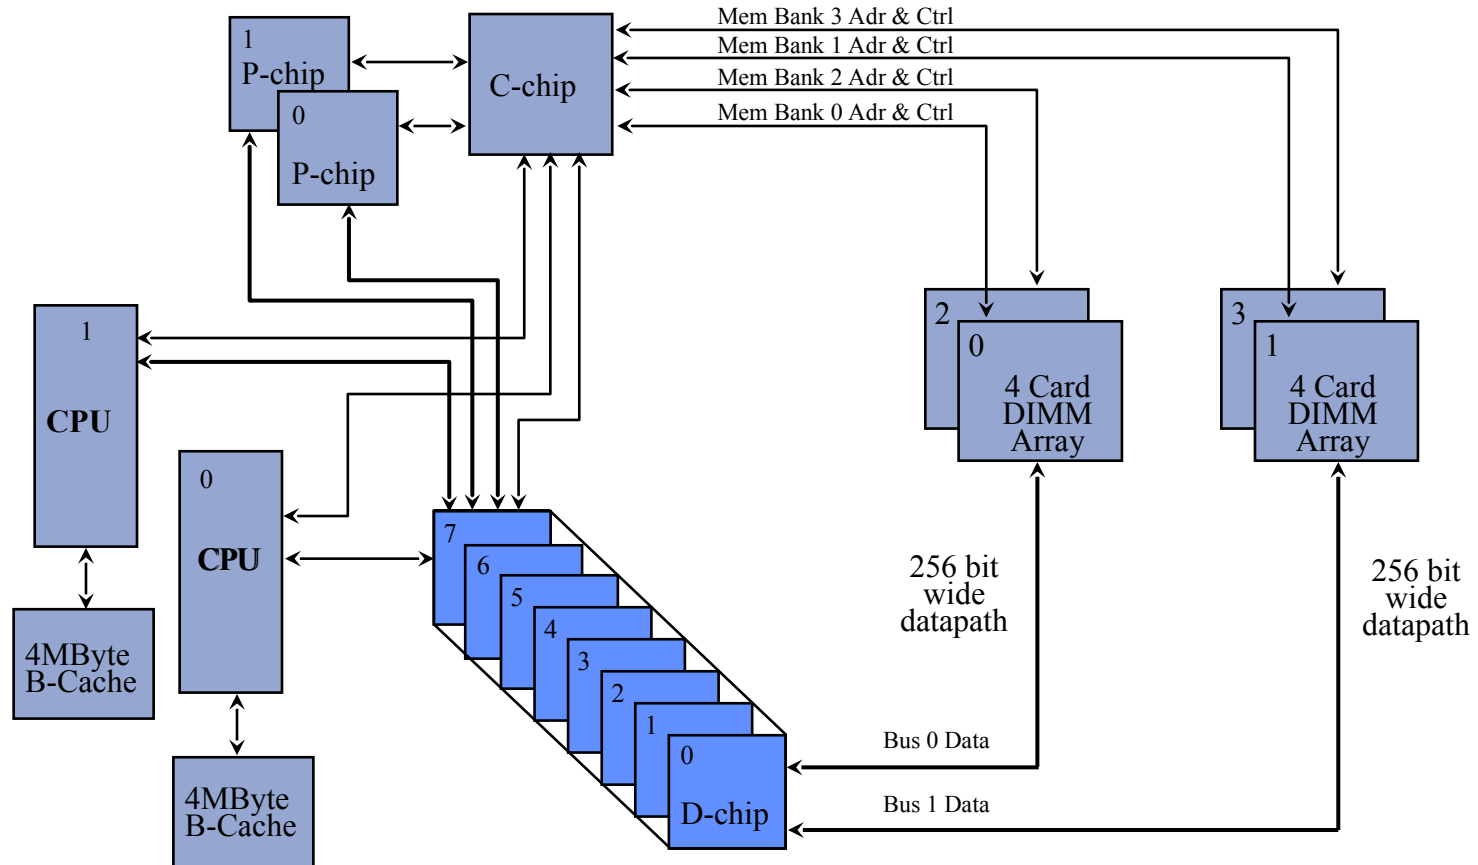
Typhoon

- C4 Chip
 - ☞ 32 GB Memory Maximum
 - ☞ 4 Memory Banks per Array
 - ☞ 4 CPU's
 - ☞ Extensive performance monitoring logic
- D Chip
 - ☞ Uses the D chip on prototypes, D4 chip in production
- ☞ MotherBoard
 - ☞ Provides an extra pipe stage to store 1 memory cycle. This will reduce the effects of late probes

Tsunami

- C2 Chip
 - ☞ 4 GB Memory Maximum
 - ☞ 2 Memory Banks per Array
 - ☞ 2 CPU's
 - ☞ None
- D Chip
 - ☞ Uses the D chip
- ☞ MotherBoard
 - ☞ None

Tsunami - Block Diagram



- Simple CPU Memory Request Sequence


CPU0 queues a read request to the Tsunami chipset. The Tsunami c-chip sends the address to CPU1 and issues the read address to the memory at the same time. The time it takes for the CPU1 to send a probe response back to the c-chip depends on where the data is in CPU1.

If the probe response is slower than the memory response, the memory data is discarded. In this case the Tsunami chipset will reissue the memory read if the probe response is a miss. If the probe response is a hit then CPU1 will supply the data to the d-chips and the d-chips will in turn supply the data to CPU0.

- Simple IO Memory Request Sequence

This same sequence is used when IO does a read except both CPU0 and CPU1 will receive a probe if IO issued the read. There is enough probe bandwidth in the Tsunami chipset to prevent probes from limiting the performance of the dual processor Goldrush except in the dropped memory data case. The dropped memory data case has been estimated to cost I/O performance of between 5 to 10%. There was no estimate for the impact on CPUs.

Tsunami - Block Diagram



Information on Goldrush and Tsunami

<http://rwhid1.eng.pko.dec.com/goldrush>