Multiprocessor Features of a PA-RISC Processor Interface Chip

**Syrus Ziai, Tom Alexander, Ken Chan,
Chuen Hu, Nazeem Noordeen, Yeffi VanAtta
Hewlett-Packard**

**Hot Chips IV
Aug 10, 1992
Stanford, California**

- *Introduction*
- System Organization
- Multiprocessor Features
- Performance Results
- Summary

Design Goals

- General Purpose Multiuser System
  - Technical and Commercial Applications
- Scalable Performance
  - Uniprocessor to > 8-way Multiprocessor
- Platform for Multiple Generations of Systems
  - Independent Upgrades of Processor, Memory and I/O

Design Approach

- Processor
  - High integer and FP throughput
  - Large first-level external cache 1-4MB FIRST LEVEL
- Memory
  - Large physical memory
  - Highly interleaved
- I/O
  - High connectivity to mass storage
  - High bandwidth bus adapter
- Bus
  - Low latency
  - High bandwidth

- Introduction
- *System Organization*
- Multiprocessor Features
- Performance Results
- Summary

- DO MEM DATA RETURN, CACHE SNOOPS HAVE TO OCCUR AT THE SAME TIME?
- IS DIRTY DATA PUSH 2x MEM ACCESS?

System Organization

- SHARED MEMORY SNOOPY.
- WIDE BUS, CUSTOM ELECTRICAL LOGIC

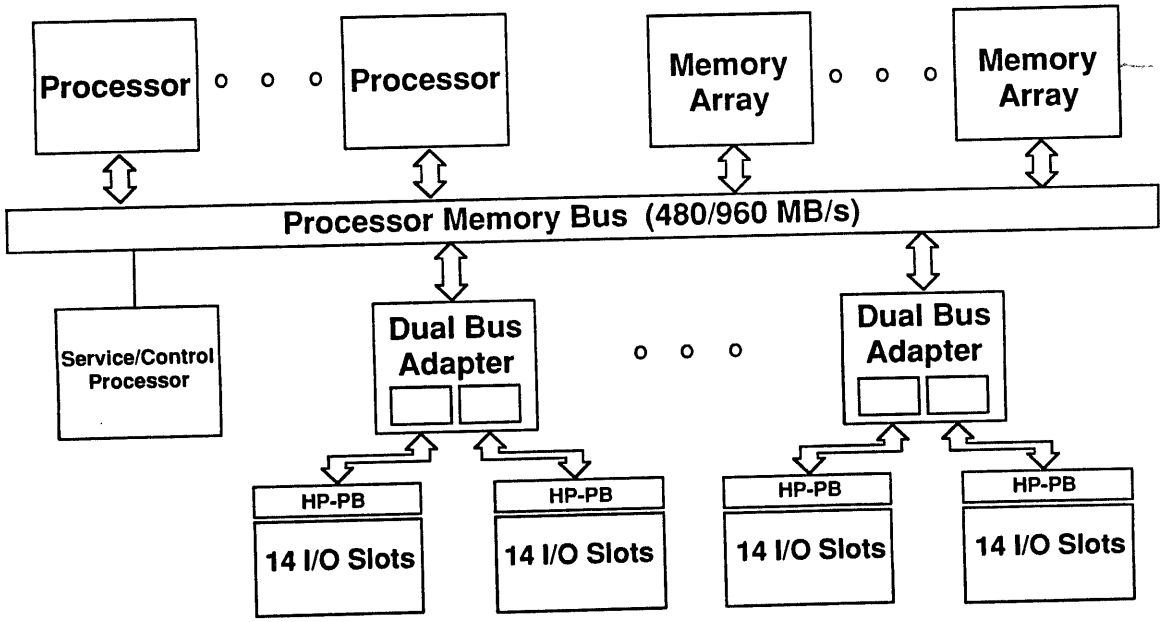

Processor Board

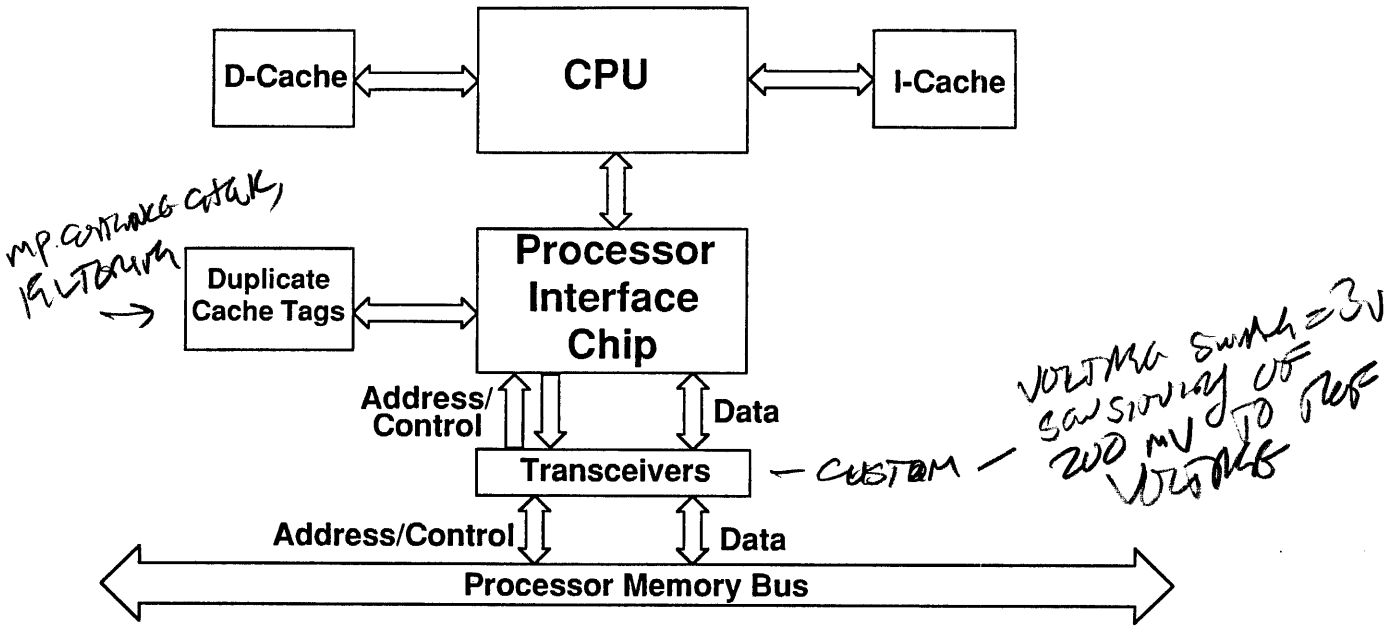

- Introduction
- System Organization
- *Multiprocessor Features*
- Performance Results
- Summary

Bus Specifications

- 16-slot backplane
- 60-MHz worst case frequency
- 960 or 480 megabyte/second bandwidth
- 128-bit or 64-bit data bus
- Separate data and address paths
- Data protected with ECC, address with parity

Synchronous Pipelined Bus

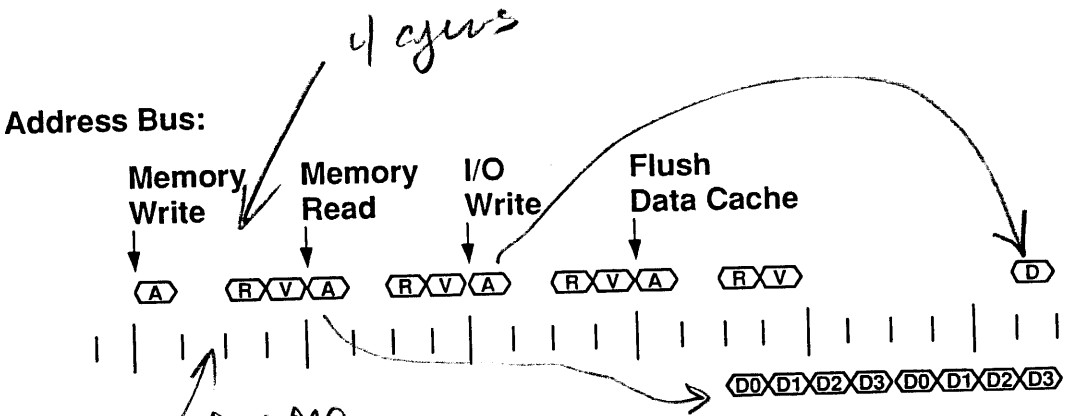

Bandwidth Enhancement Features

- Flexible memory interleaving
  - Up to 64-way interleaving for each of 3 groups
  - * Unified memory interleaving among non power of 2 modules
- Resource-driven arbitration
  - Memory block status maintained by masters
  - Transactions not queued by memory
- Slave-driven arbitration
  - I/O re-arbitrates when data available

Cache Coherence Transaction

- Bus protocol for cache and TLB coherence
- Processor-to-processor data transfer

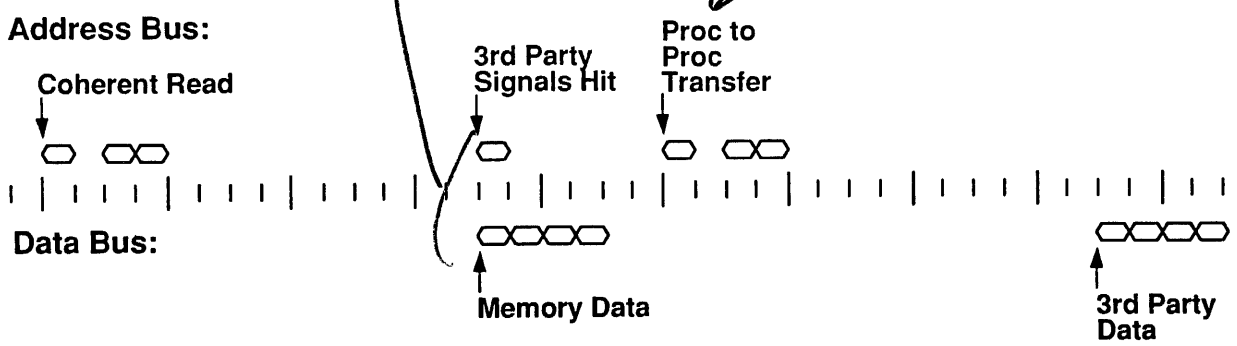

Master Processor
3rd Party Processor
Memory

Cache Coherence (cont'd)

- Snoopy cache coherence with duplicate cache tags
- Multiple outstanding coherence checks

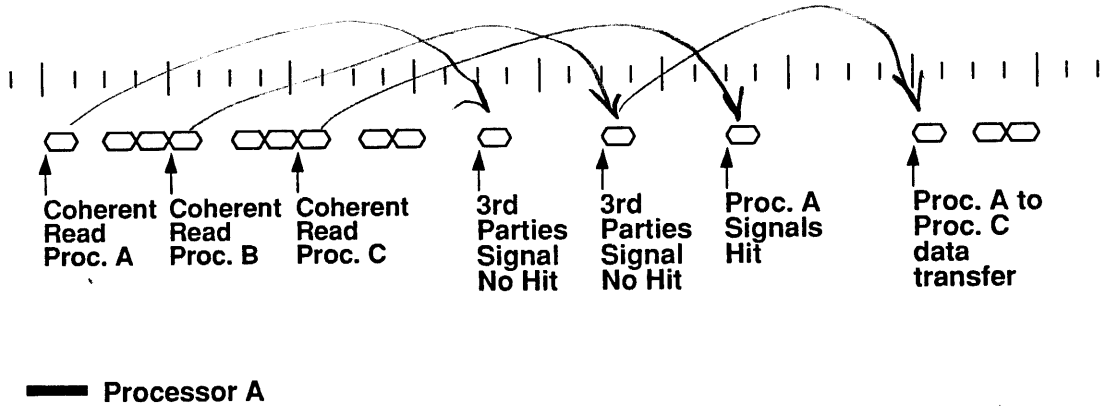

Write Buffers

- Circular FIFO with random access

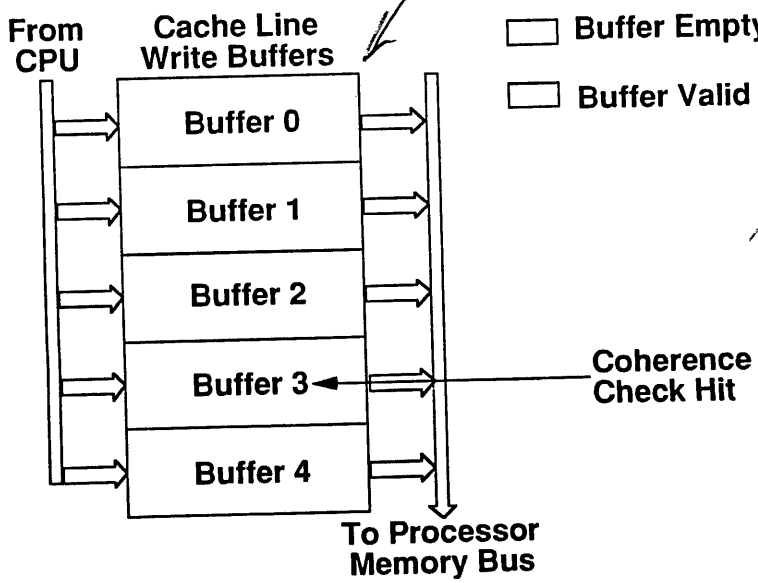

Multiprocessor Tuning Features

- **Configurable bus parameters**
  - **Memory latency and cycle time**
  - **Coherence check**
  - **Interleave factor**
- **Pipeline freeze capability**
  - **Coherence exception conditions**
- **Performance monitoring counters**
  - **Multiprocessor events**
- **Support multiple generations of CPUs**

die photograph

- Introduction
- System Organization
- Multiprocessor Features
- *Performance Results*
- Summary

90% of 1.7 SOC

TPCA - 578 TPS (4-way, 60MHz CPU 60MHz BUS)

SPEC	RATE	1174	2360	3529	4685
	FP 92	(1.0)	(2.01)	(3.0)	(3.99)
SPEC	RATE INT 92	1164	2253	3306	4301
		(1.0)	(1.94)	2.84)	(3.70)

Only (Compress, GCC don't scale well due to disk activity)

Performance Results

For PA 1.1 SC 1.0 (not 7120)

This slide will present UP and MP performance results.

Summary

Demonstrated a highly scalable many-way MP system using high performance RISC CPUs

Achieved 8-way MP with first silicon

DREW QA OF 9 MONTHS.