


The Orca Chip . . .
Heart of IBM's RISC System/6000
“Value” Servers

Ravi Arimilli
IBM RISC System/6000 Division

- **Server Background**
- **Cache Heirarchy Performance Study**
- **RS/6000 Value Server System Structure**
- **Orca Overview**
- **Orca Design Points**
- **System Comparisons**
- **Summary**

- 2H 94 IBM Announced the RISC System/6000 Model J30/R30:



- **IBM's First PowerPC SMP System**
- **1–8 Way SMP Scalable Structure**
- **Introduction of the PowerPC 6XX System Bus**
- **Symmetric and Coherent I/O Subsystems**
- **New and Efficient OS for SMP (AIX Ver. 4)**
- **Upgradable Processor Cards**
 - **Higher Frequency Processors**
 - **New PowerPC Processor Designs (604/604e/620/etc)**
- **“Baseline” Platform to Begin Performance Analysis**

- **TPC-C**
 - **Mix of database transactions w/high multiprogramming level**
 - **Multiple implementations using different database products**
 - **Large code and data footprint**
 - **Not OS intensive (12% OS, 88% DB)**

- **SPEC SFS 097.laddis**
 - **NFS file server benchmark**
 - **Mix of file server transactions w/high multiprogramming level**
 - **Modest code and data footprint**
 - **Mostly OS intensive (100% OS)**

- **SPEC SDM 057.sdet**
 - **Batch oriented software development benchmark**
 - **Modest multiprogramming level**
 - **Large code and data footprint**
 - **Mostly OS intensive (50% OS, 50% commands/libraries)**

- **SPEC SDM 061.kenbus**
 - **Interactive oriented multiuser benchmark**
 - **High multiprogramming level**
 - **Large code and data footprint**
 - **Mostly OS intensive (50% OS, 50% commands/libraries)**

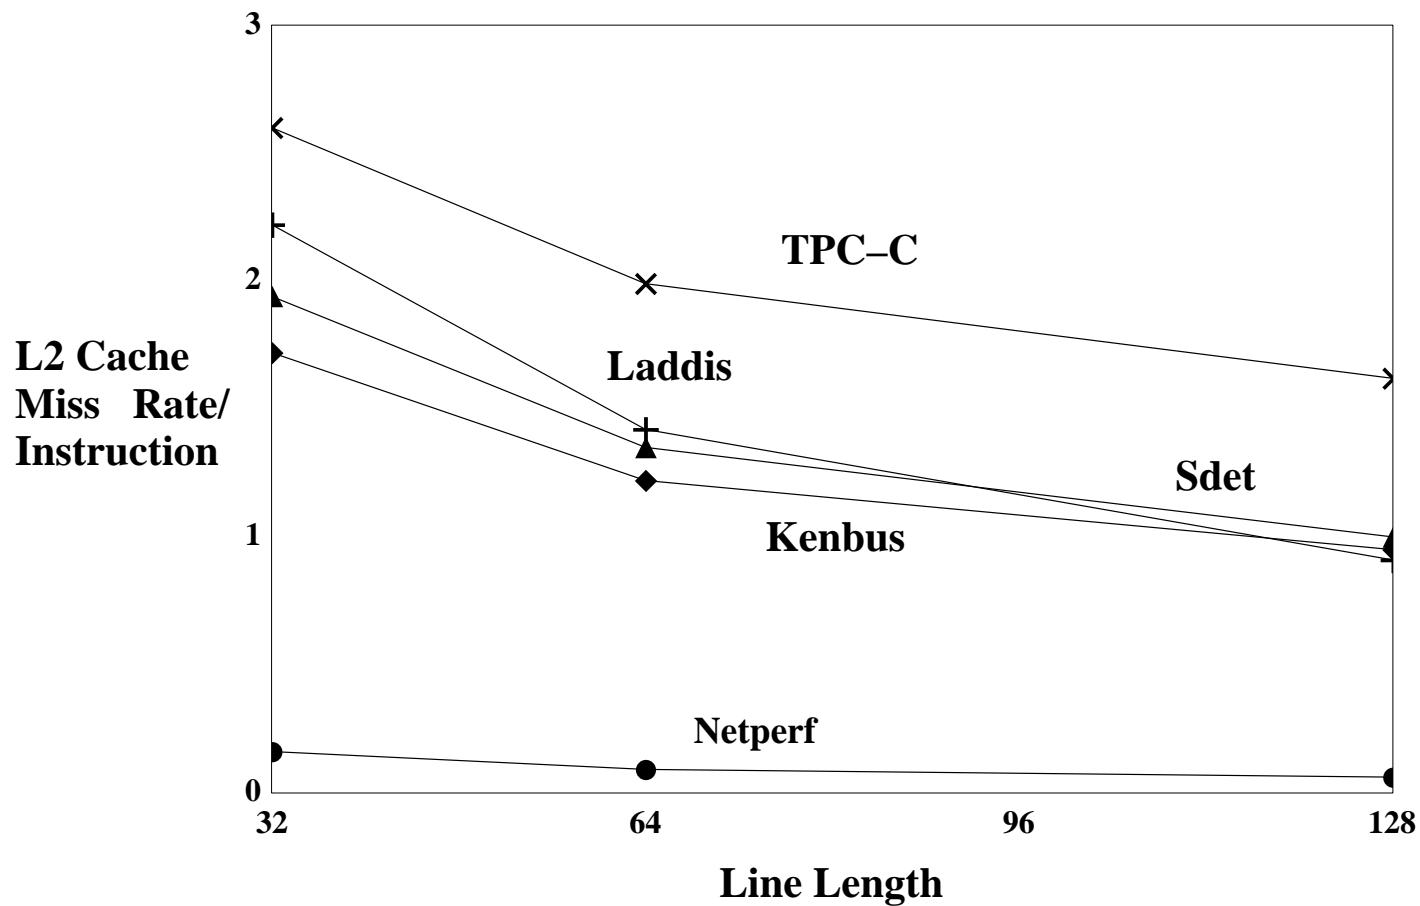
- **Netperf**
 - **TCP-IP Performance Test**
 - **Mostly sends/receives variable data lengths in loop format**
 - **Small code footprint, mostly runs in L1 Cache**
 - **Sensitive to processor Mhz**

- **Other**
 - **SPECint 95**
 - **SPECfp 95**
 - **G92 (Computational Chemistry)**
 - **Les (Aircraft surface turbulence simulation)**

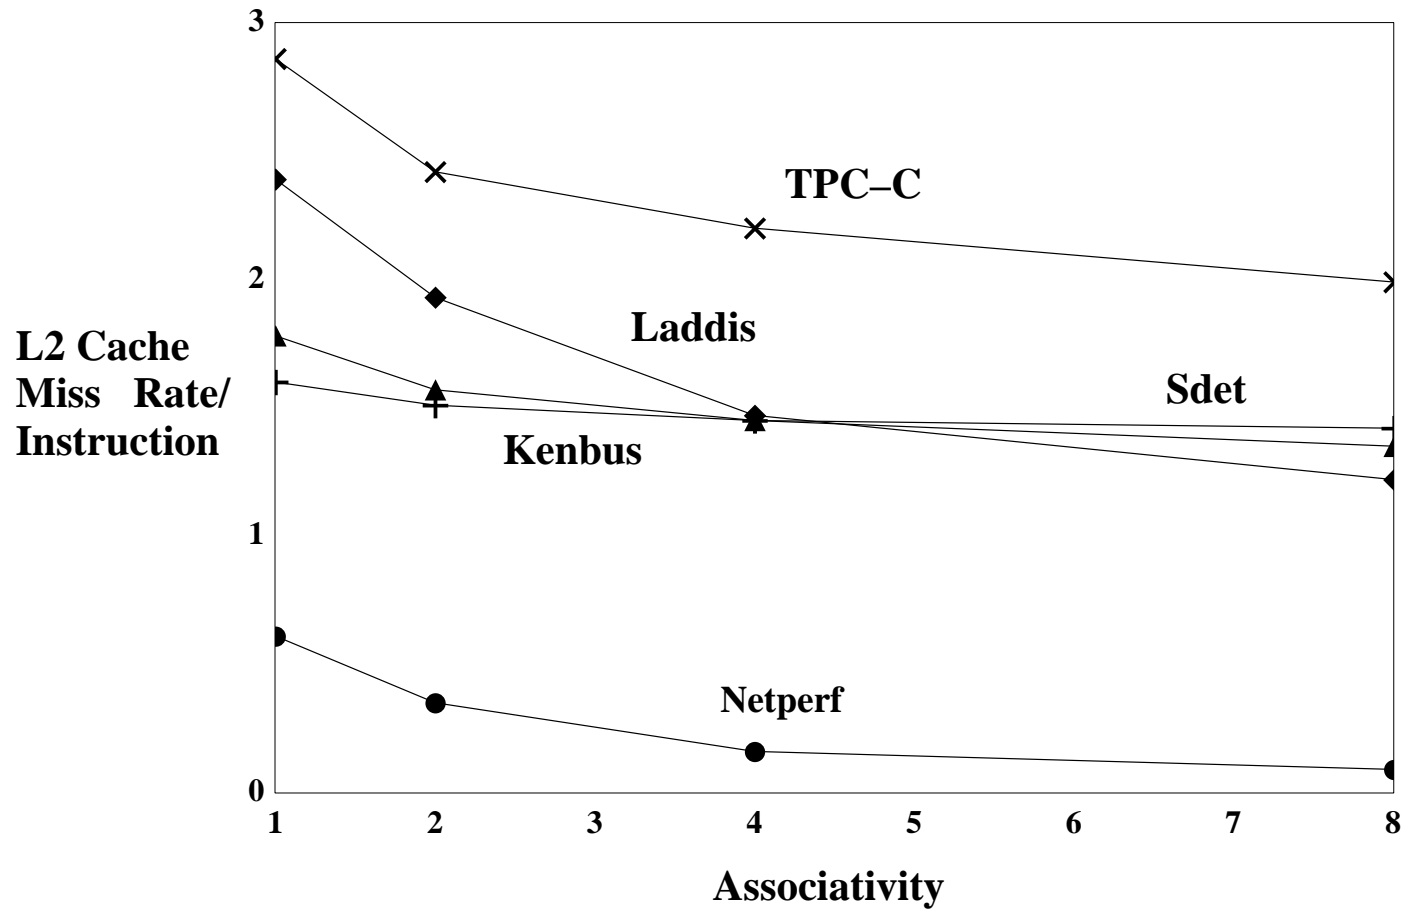
- **Performance Tools**
 - **Software Instruction Trace Tools**
 - **Hardware Bus Trace Tools**
 - **Processor Simulators**
 - **Cache Simulators**
 - **Memory Simulators**
 - **System Topology/Interconnect Behaviorals**
 - **Validation Tools**

- **L2 Cache Parameters Varied**
 - **Processor to L2 Bus Ratios (1:1, 3:2, 2:1, etc)**
 - **Associativity**
 - **Cache Line Size**
 - **Sectors/Cache Line**
 - **Cache Replacement Algorithms**
 - **L2 Access Latency**
 - **L2 Intervention Latency**
 - **Lookaside vs Inline L2 Caches**
 - **Unified vs Split I/D Caches**
 - **Shared L2 Cache vs Dedicated L2 Cache**

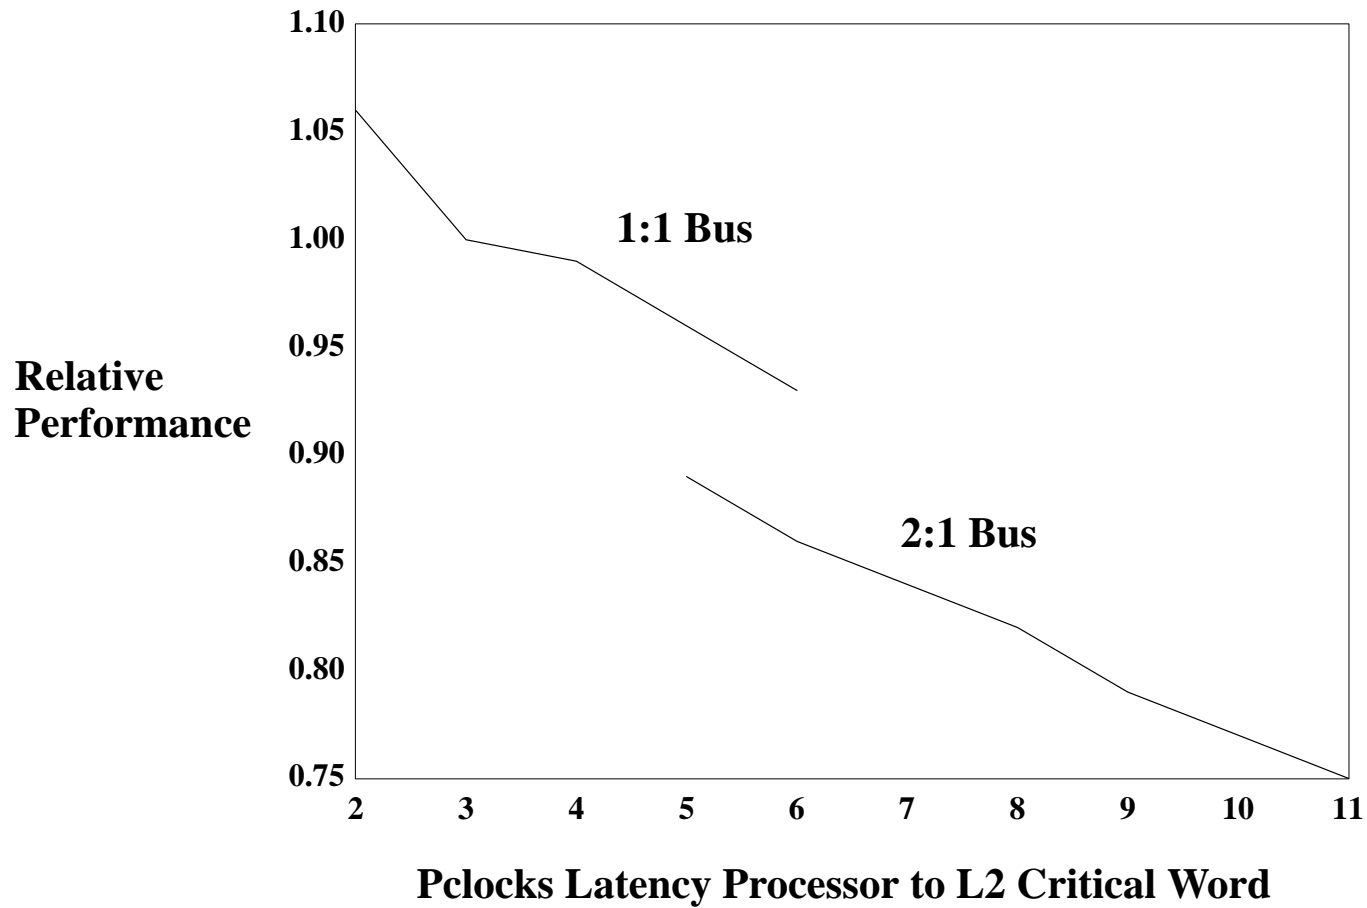
- **System Parameters Varied**
 - **601, 604, 604e, and 620 Processor Cores**
 - **Memory Access Latency**
 - **System Bus Width**
 - **System Bus Ratios**
 - **Memory Bus Width**
 - **Memory Bus Ratios**
 - **Switch vs Bus for System Address Bus and Data Bus**
 - **Switch vs Bus for Memory Bus**
 - **Number of Processors**
 - **System Bus Protocols**



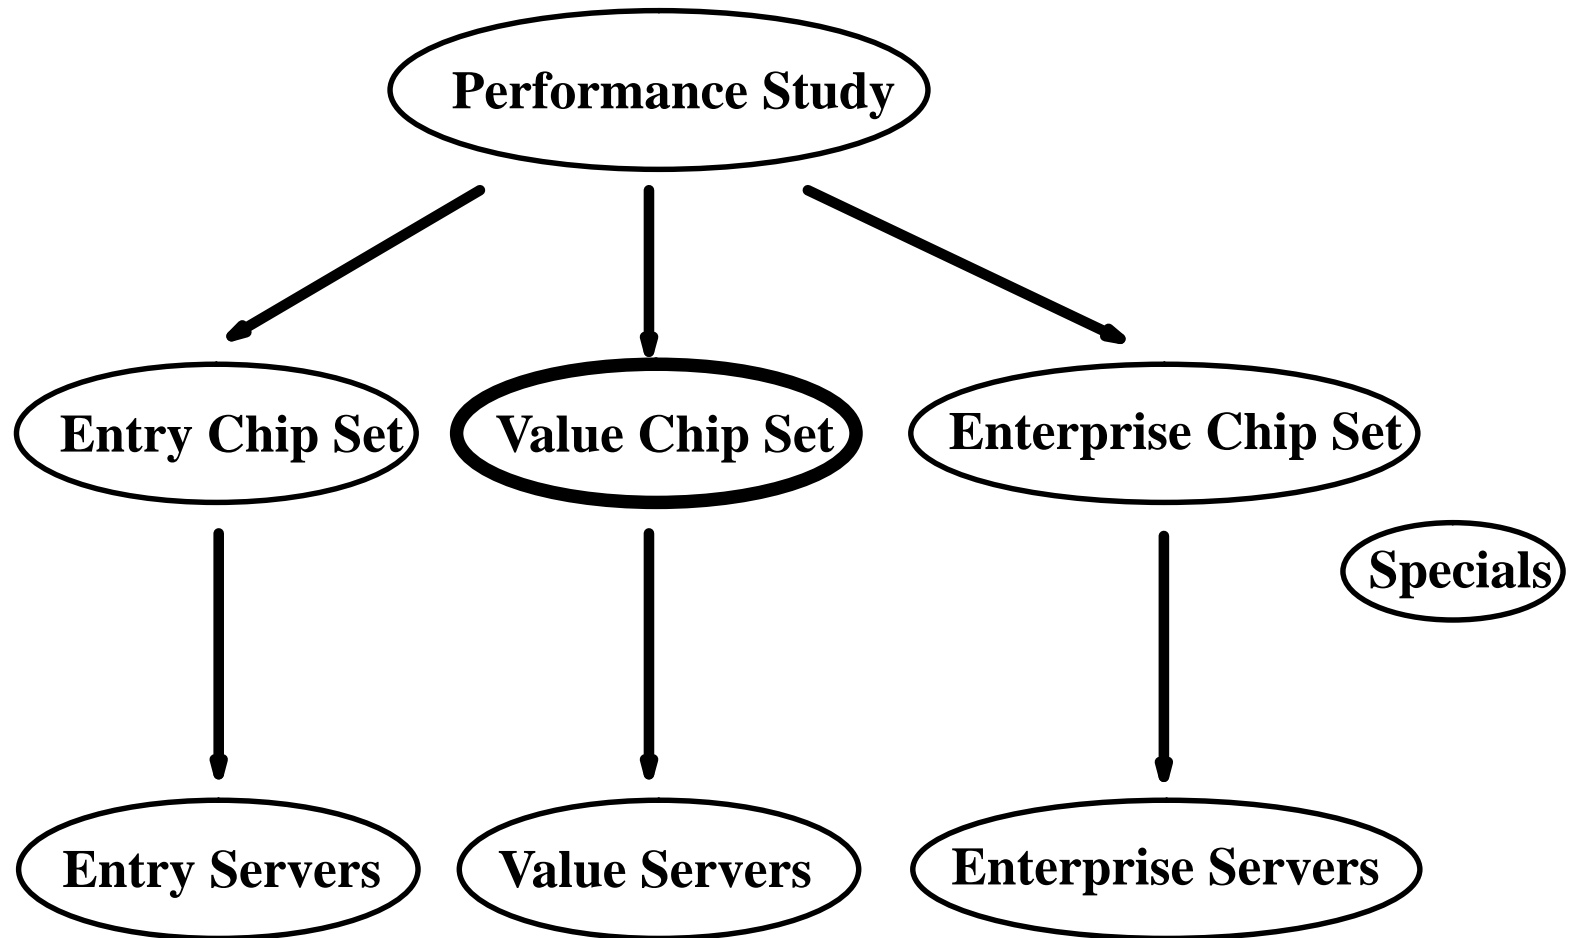
Effect of Increasing Line Length on Miss Rates (256KB 8-Way SA)




Effects of Increasing Associativity on Miss Rates (256KB 64B Line)



Effect of L2 Latency on Performance of TPC-C

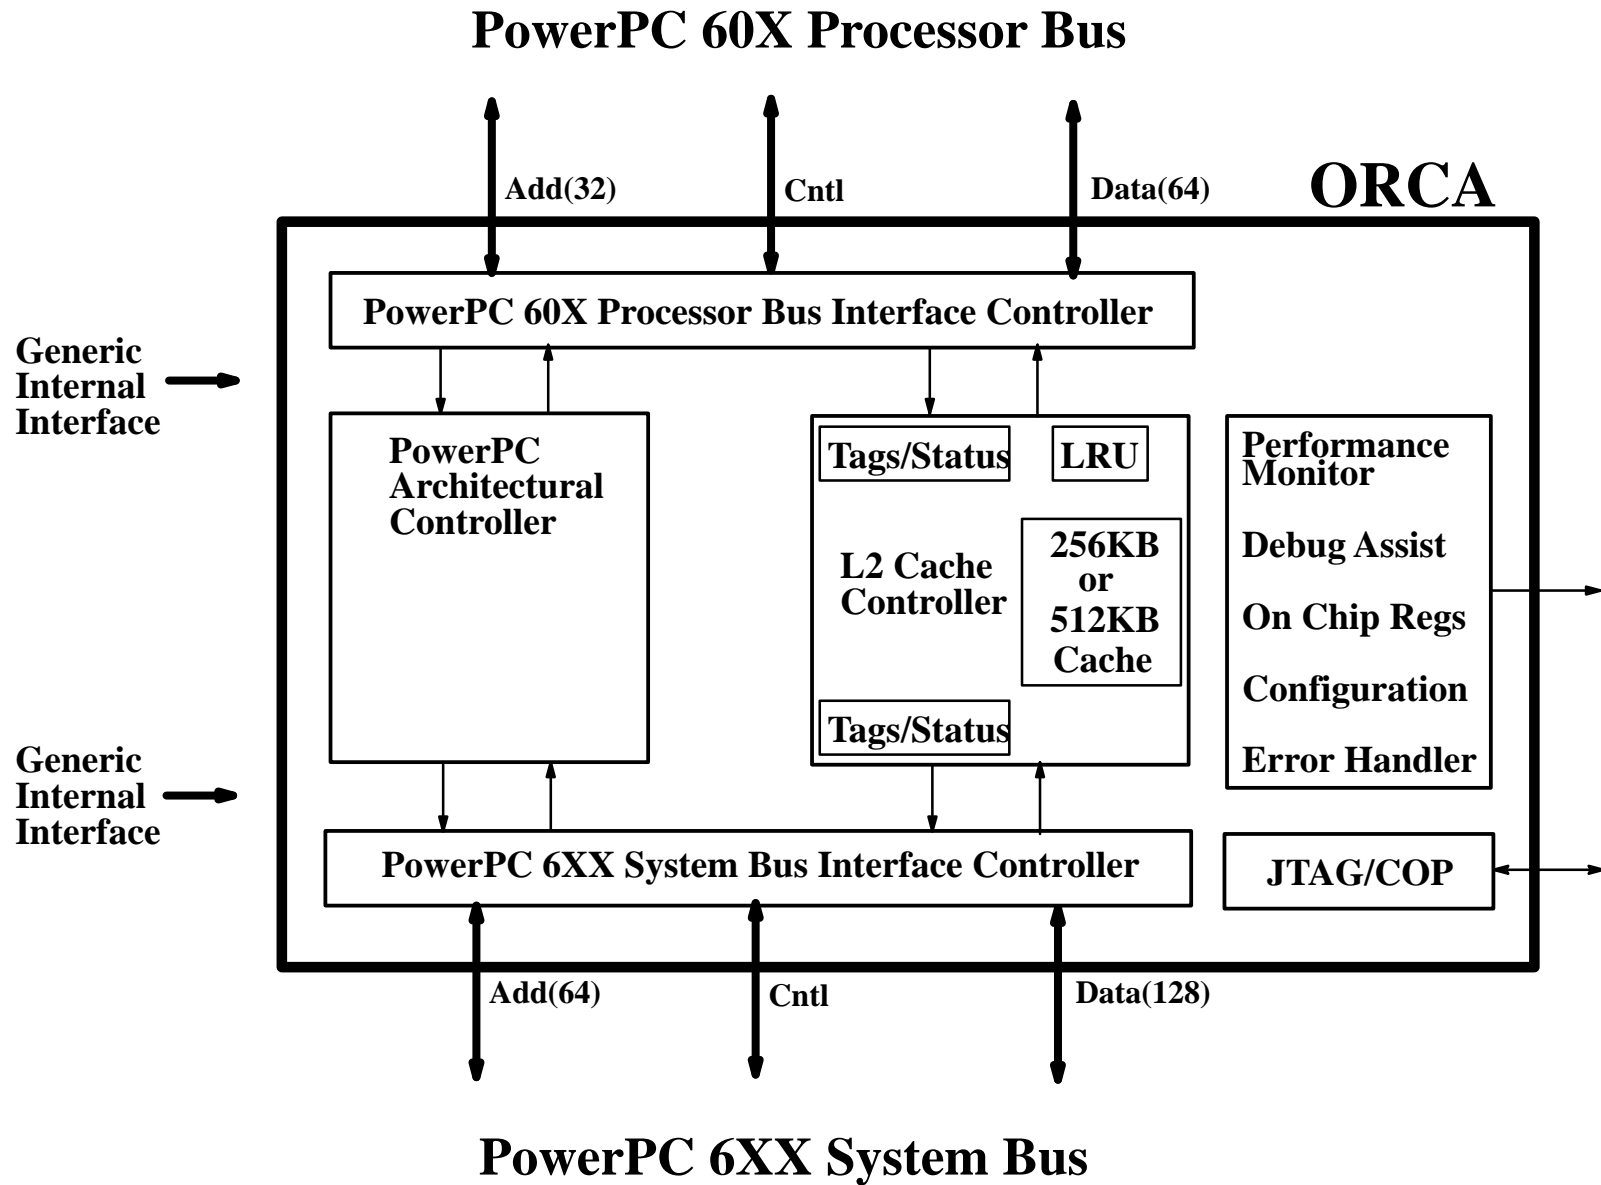


RS/6000 Value Server Structure



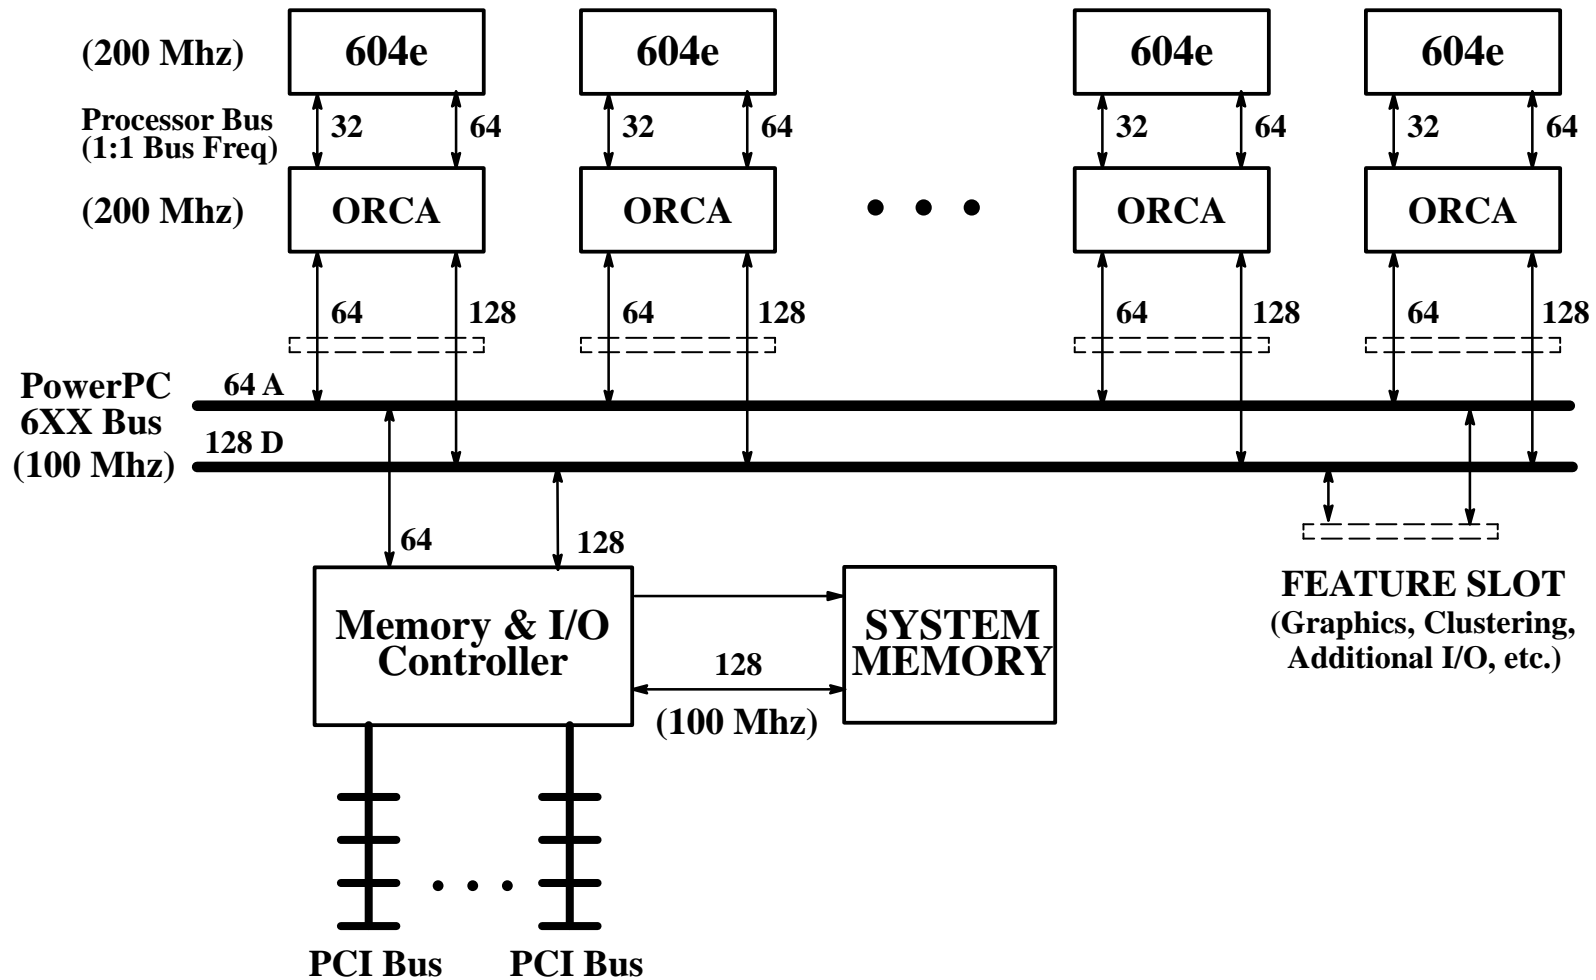
- Fully integrated Cache, Tags, Status, LRU and Controller**
- Cache Controller Features**
 - Fully Integrated L2 Cache, Directory, and Controller**
 - 256KB or 512KB Cache Sizes**
 - 8 Way Set Associative**
 - Low Latency L1 Miss, L2 Hit Access**
 - 200 Mhz Operation (1:1 with CPU Frequency)**
 - 32B/Cycle Internal Cache Access (6.4 GB/sec Cache Bandwidth)**
 - 64B L2 Cache Lines (32B L1 Cache Lines)**
 - Non-Blocking L2 Cache for Processor Bus Reads and Writes**
 - Non-Blocking L2 Directory for Processor Bus Snoops**
 - Non-Blocking L2 Directory for System Bus Snoops**
 - Single Cycle Snoop Coherency Resolution**
 - High Speed System Bus Intervention Supported**
 - Queue Depths Support Processor Bus and System Bus Saturation Limits (Improves Technical/fp Performance)**
 - RISC Style Micro-Architecture (Shallow Logic/Heavily Pipelined)**
 - Improved SMP “Locks” Performance**
 - High Speed Cache Inhibited Stores (Graphics Performance)**
 - Single Bit Error Recovery (ECC) for Internal Cache & Directory**

- **PowerPC 6XX System Bus Features**
 - **Seperate Address and Data Buses (Fully Tagged)**
 - **Efficient Address and Data Bus Arbitration Protocols**
 - **64 Bit Addressing Support**
 - **64 Bit and 128 Bit Data Bus Support**
 - **High Speed Intervention Support**
 - **Split “Flow Control” and “Coherency” Responses**
 - **Efficient Protocols for SMP Clustering**
 - **Efficient Protocols for Switched Address and Data Buses**
 - **High Frequency Capable (Latch to Latch Protocols, etc)**
 - **“System Centric” Bus Architecture (The “System” Directly Controls All OCD Enables, When Snoopers Sample, etc)**
 - **Heavily Pipelined Address/Response Buses**
 - **Other (Bus Parking, Prefetch Hints, Large Bursts, etc)**
 - **Robust Error Recovery**

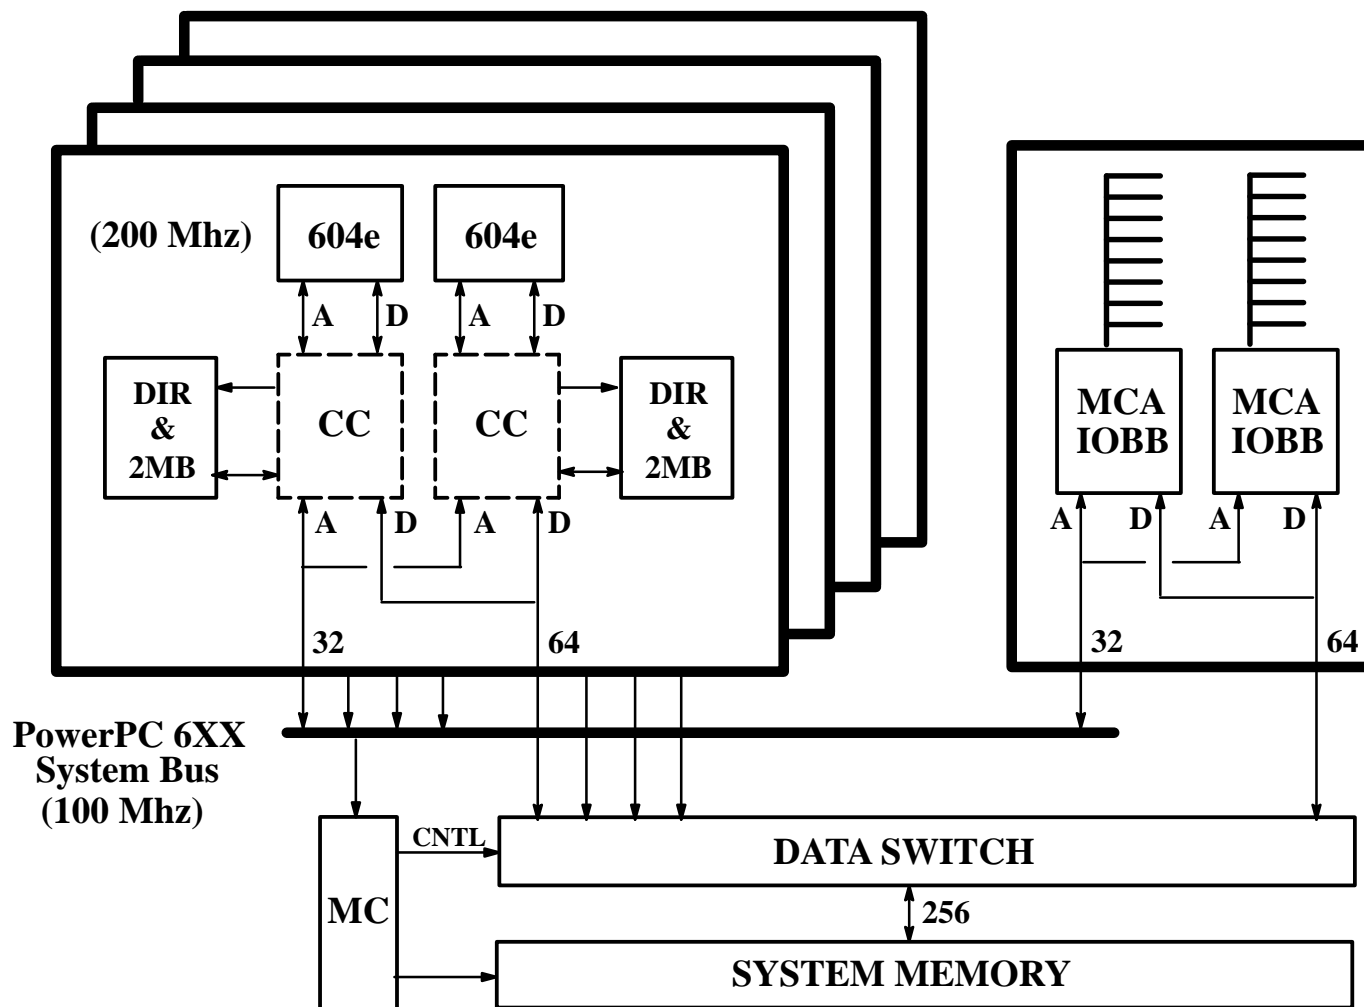


- **256KB, 512KB L2 Cache**
- **0.5u Nwell CMOS**
- **Five Level Metal**
- **Local Interconnect**
- **$L_{\text{eff}} = 0.25\mu\text{m}$**
- **2.5V Core Voltage**
- **3.3V Drivers/Receivers**
- **7.5W Typical Power at 200 Mhz (estimated)**
- **430 I/O Signals, CMOS/TTL Compatible**
- **LSSD Design, JTAG Compliant**
- **32mm Ball Grid Array**

RS/6000 Value Server Structure



- **Alternative Option via RS/6000 J30/R30 Parts**



<u>Parameters</u>	<u>System w/R30 Parts</u>	<u>System w/Value Parts</u>
Processor/Frequency	604e @ 200Mhz	604e @ 200Mhz
Processor I/D, Associativity	32K/32K, 4Way	32K/32K, 4Way
Processor Bus Frequency	100 Mhz (2:1)	200 Mhz (1:1)
L2 Access Latency (Pclks)	9-2-2-2	3-1-1-1
L2 Cache Size	512K, 1MB, 2MB	256K, 512K
L2 Associativity	1 Way	8 Way
L2 Cache Line Size	32 Byte	64 Byte
L2 Sectors/Cache Line	1 Sector	2 Sectors
L2 Inline or Lookaside	Inline	Inline
L1 Inclusivity	Imprecise	Precise
L2 Unified or Split I/D	Unified	Unified
L2 Shared or Dedicated	Dedicated	Dedicated
SB Address Switch or Bus	Bus	Bus
SB Data Switch or Bus	Switch	Bus
SB Data Bus Width	8 Bytes	16 Bytes
SB Intervention Latency	Slow	Fast
Memory Bus Width	32 Bytes	16 Bytes
Memory Bus Frequency	50 Mhz	100 Mhz
I/O Subsystem	Micro Channel	PCI

- **IBM has performed extensive UNIX server performance studies.**
- **The result of these studies has led to the development of three server chip sets within the IBM RS/6000 Division:**
 - **Entry Servers**
 - **Value Servers**
 - **Enterprise Servers**
- **The heart of the Value Servers is the Orca Chip:**
 - **Fully Integrated L2 Cache (256KB/512KB), Directory, & Controller**
 - **Drastic departure from traditional L2 Cache Controller designs**
 - **Highly associative (8 way), low latency, high bandwidth design**
 - **Aggressive, fully non-blocking, heavily pipelined design points**
 - **Initial offering at 200Mhz with future frequency increases**
 - **Supports the 128 bit PowerPC 6XX System Bus**
 - **Robust Performance Monitor Support**
- **The ORCA Chip provides high commercial performance, and scalability w.r.t. the “number” and “frequency” of processors.**
- **The ORCA Chip enables cost effective desktop PowerPC Microprocessors (604e) to be used in SMP Server environments.**