



A Multi-threaded 64 Bit PowerPC Commercial RISC Processor

IBM Server Group
Rochester, MN



Authors

Salvatore N. Storino
John M. Borkenhagen
Ronald N. Kalla
Steven R. Kunkel

Email: storino@us.ibm.com



Goals

- Very High Performance on Commercial* Benchmarks
- Scalability up to 24-Way Processors
- Multi-threading Operation up to Virtual 48-Way
- High L2 and Memory Bandwidth
- Low Latency Design
- Optimized Power, Area and Performance

* Commercial = TPC, SAP, SpecWeb and Lotus Notes



Outline

- | | |
|------------------------|--|
| 1. Commercial Overview | 8. Peak Performance on Memcpy |
| 2. CPI Breakdown | 9. Architecture for MT control registers |
| 3. L2 Miss | 10. MPs Configuration |
| 4. Processor Overview | 11. Performance Improvements with MPs |
| 5. Pipeline | 12. Industry Standard Benchmarks |
| 6. MT Structures | 13. Summary |
| 7. MT GPR Cell | 14. Chip Plot |



Commercial Server Code

Typical String:

- 1) LD; Fetch Record
- 2) CMP; Match Key
- 3) BC; Branch to Handler

Characteristics

Seldom Loops, as in Matrix Algebra
Large Working Set, frequent Read-Write sharing
Branch Prediction Difficult



Characteristics of Commercial Server Workloads

Natural Task Level Parallelism

Large Number of Users
Throughput is Primary Performance Measurement

Significant Cache and Memory Component to CPI

Trends in Memory Latency are Increasing

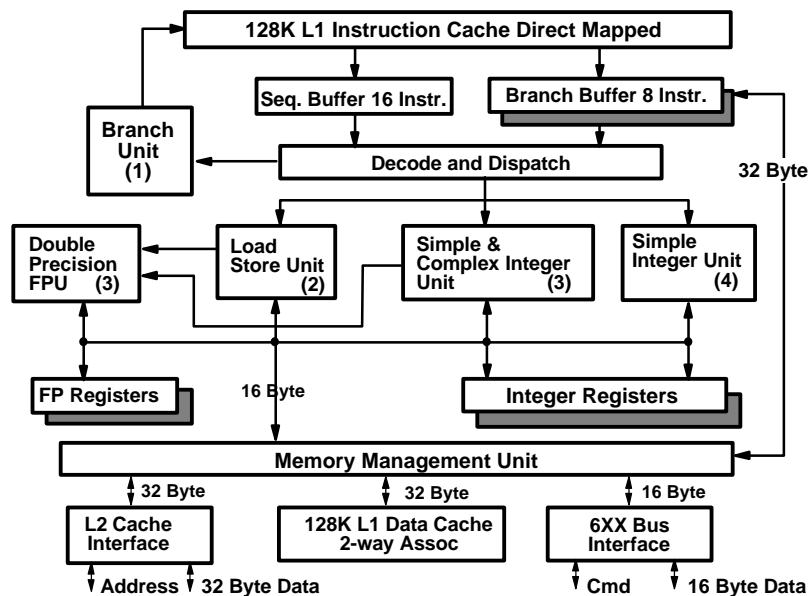


Processor Migration

	<u>Northstar</u>	<u>Pulsar</u>
Frequency	350 MHz	450 MHz
MPs	12 Way	24 Way
Cache Size	128 KB	128 KB
L2 Peak BW	8.4 GB/s	14.4 GB/s
Mem Peak BW	1.9 GB/s	2.4 GB/s
Mem Bus Ratio	3:1	3:1
Max Memory	40 GB	96 GB
Chip Size	162 sqmm	139 sqmm
Power	34W	22W
Transistors	12M	34M
Technology	6S2, 2.5v	7S, 1.8v
Metallurgy	5-Al Bulk	6-Cu Bulk

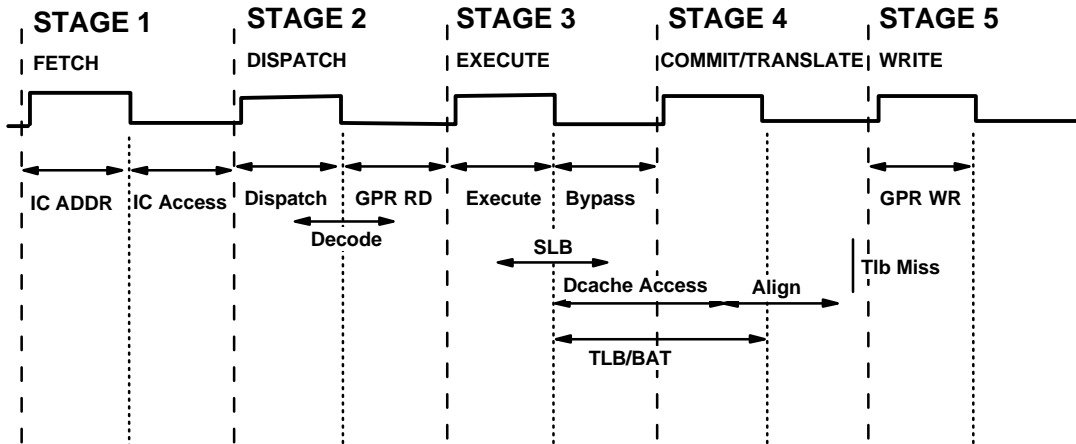


Processor Overview

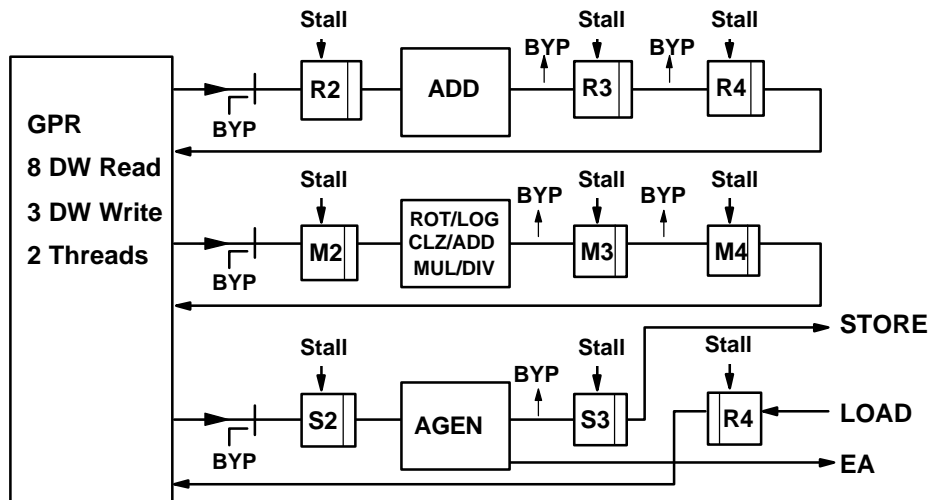




Star Series PipeLine



FX Architectural View





Multi-Threaded Structures

MT Structures

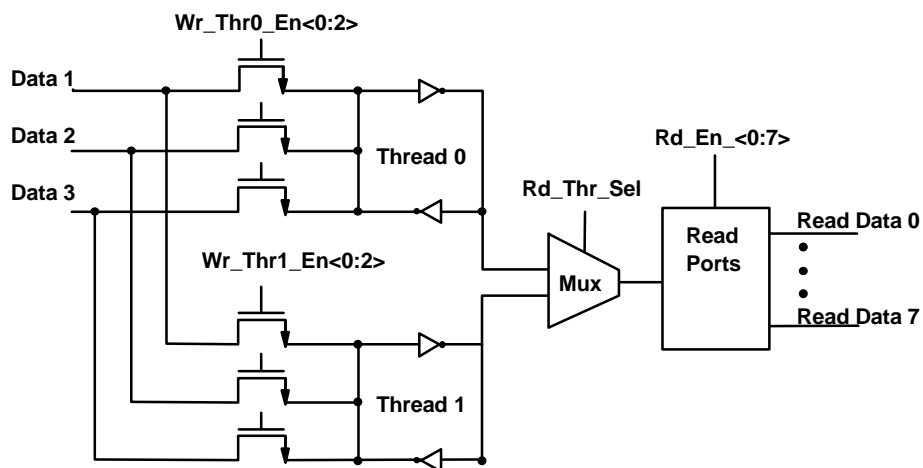
- GPR, FPR and SPRs
- Queue Expansion for Extra Traffic

Non MT Structures

- IC, DC, TLB, FUs and MMU




Multi-Threaded GPR Cell






"Memcpy" Diagram




Source or Destination Miss



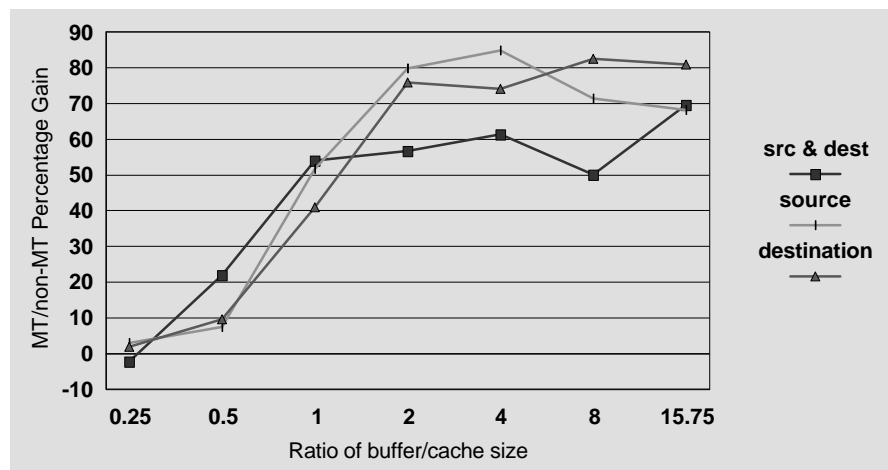
non-MT Execution



MT Execution



Pulsar Performance Improvement for Memcpy applying MT



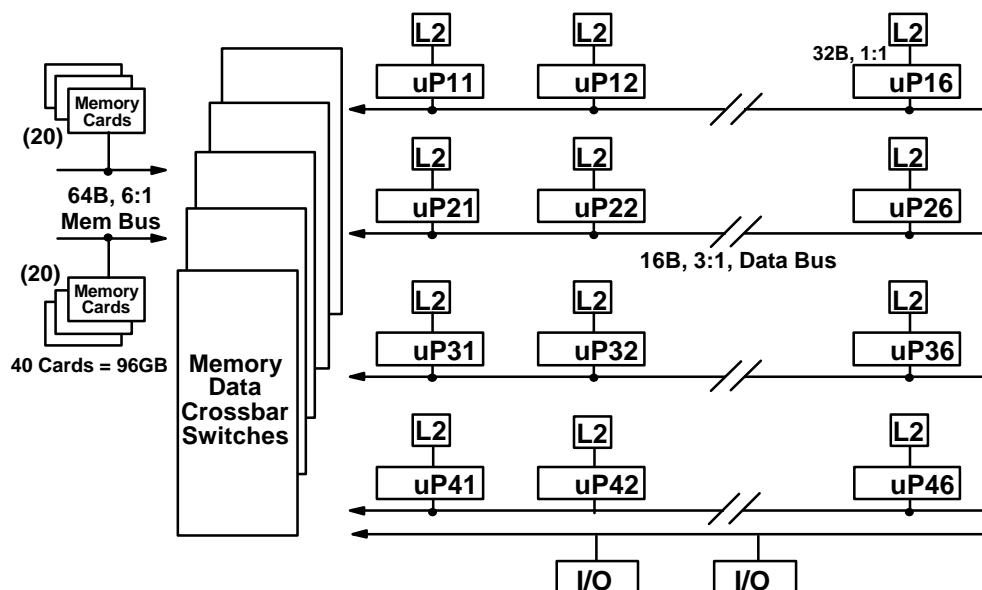


Architected MT Control Registers

- TSC<0:31> Thread Switch Control
 - Controls events for thread switching
 - L1 misses, L2 misses, TLB misses and etc.
- Priority Control Instructions
 - NOPs that set thread priority to lo, med, hi
 - NOPs are backward compatible
- TST<22:31> Thread Switch Time-out
 - Limits cycles that one thread is active
 - prevents "infinite loop" task domination
- Forward Process Guarantee
 - limits number of thread swaps from an instruction
 - prevents threads from thrashing




Pulsar MP 24-Way System

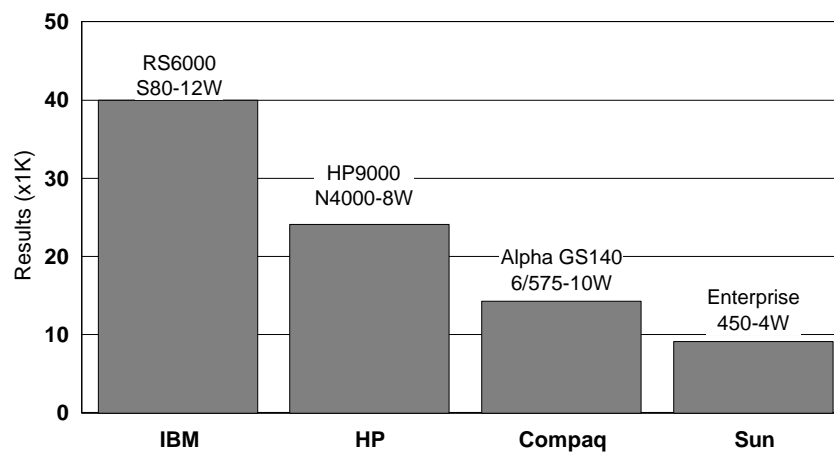




Pulsar Performance Improvement MT versus MPs



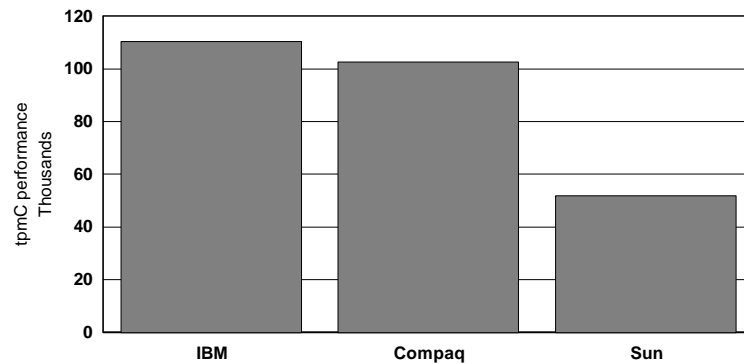
Top SPECweb Performance



Data as of: 13 July 99



Top TPC-C Clustered Systems



IBM	Compaq	Sun
RS6000	Alpha	Enterprise
S7A	8400	6000
8nodex12w	8nodex12w	2nodex22w
\$122.44	\$139.49	\$134.46
\$/tpmC	\$/tpmC	\$/tpmC

* Data as of 07/13/99 from:
<http://www.ideasinternational.com/benchmark/bench.html>



Follow On Direction

- More MPs
- NUMA Options
- Bigger L2
- SMI interface
- Premier Technology "SOI"
- Low E Wire Technology
- Frequency Growth
- Constant Power Curve