Case Study of a Multi-core, Multi-threaded Processor—The Sun T1 (“Niagara”), T2, T3

55:132/22C:182

Reference

- **Computer Architecture, A Quantitative Approach, Fourth Edition**, by John Hennessy and David Patterson, chapter 4.

T1 (“Niagara”)

- Target: Commercial server applications
  - **High thread level parallelism (TLP)**
    - Large numbers of parallel client requests
  - **Low instruction level parallelism (ILP)**
    - High cache miss rates
    - Many unpredictable branches
    - Frequent load-load dependencies
- Power, cooling, and space are major concerns for data centers
- Metric: Performance/Watt/Sq. Ft.
- Approach: Multicore, Fine-grain multithreading, Simple pipeline, Small L1 caches, Shared L2

T1 Architecture

- Also ships with 6 or 4 processors

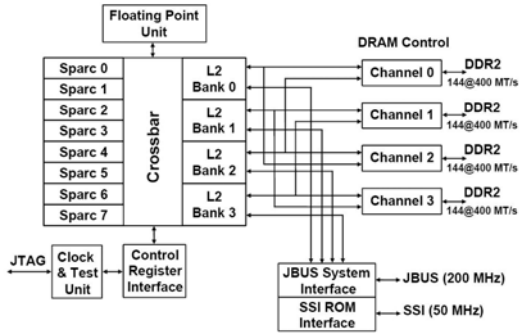

T1 pipeline

- Single issue, in-order, 6-deep pipeline: F, S, D, E, M, W
- 3 clock delays for loads & branches.
- Shared units:
  - L1 $, L2 $
  - TLB
  - X units
  - pipe registers
- Hazards:
  - Data
  - Structural

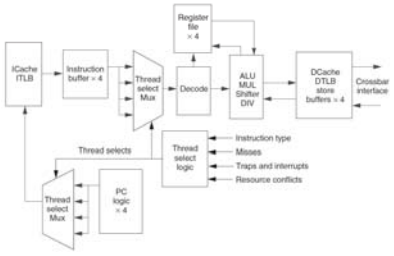

Integer Register File

- One register file / thread
- SPARC window: in, out, local registers
- Highly integrated cell structure to support 4 threads:
  - 8 windows of 32 locations / thread
  - 3 read ports + 2 write ports
  - Read/write: single cycle latency
- 1 Active Window Cell (copy of the architectural set window)

Thread Scheduling

- Thread selection based on:
  - Previous long latency instruction in pipe
  - Instruction type
  - LRU status
- Select & Fetch coupled

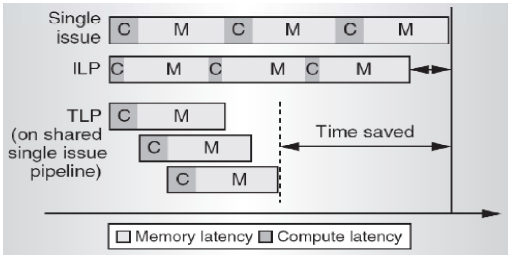

T1 Fine-Grained Multithreading

- Each core supports four threads and has its own level one caches (16KB for instructions and 8 KB for data)
- Switching to a new thread on each clock cycle
- Idle threads are bypassed in the scheduling
  - Waiting due to a pipeline delay or cache miss
  - Processor is idle only when all 4 threads are idle or stalled
- Both loads and branches incur a 3 cycle delay that can only be hidden by other threads
- A single set of floating point functional units is shared by all 8 cores
  - floating point performance was not a focus for T1

Memory, Clock, Power

- 16 KB 4 way set assoc. I$/ core
- 8 KB 4 way set assoc. D$/ core
- 3MB 12 way set assoc. L2 $ shared
  - 4 x 750KB independent banks
  - crossbar switch to connect
  - 2 cycle throughput, 8 cycle latency
  - Direct link to DRAM & Jbus
  - Manages cache coherence for the 8 cores
  - CAM based directory

  Write through
  - allocate LD
  - no-allocate ST
- Coherency is enforced among the L1 caches by a directory associated with each L2 cache block
- Used to track which L1 caches have copies of an L2 block
- By associating each L2 with a particular memory bank and enforcing the subset property, T1 can place the directory at L2 rather than at the memory, which reduces the directory overhead
- L1 data cache is write-through, only invalidation messages are required; the data can always be retrieved from the L2 cache
- 1.2 GHz at ≈72W typical, 79W peak power consumption

Miss Rates: L2 Cache Size, Block Size

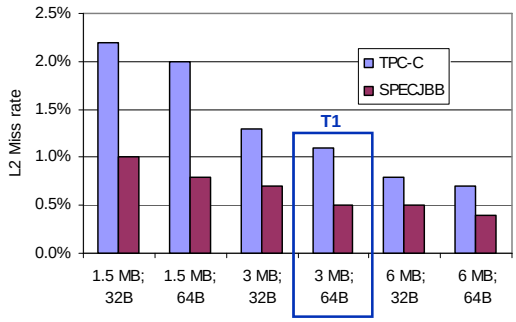

Miss Latency: L2 Cache Size, Block Size

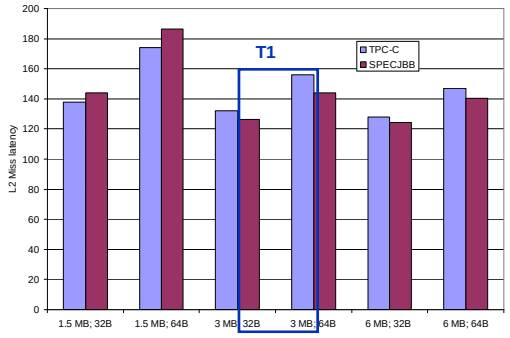

CPI Breakdown of Performance

Benchmark	Per Thread CPI	Per core CPI	Effective CPI for 8 cores	Effective IPC for 8 cores
TPC-C	7.20	1.80	0.23	4.4
SPECJBB	5.60	1.40	0.18	5.7
SPECWeb99	6.60	1.65	0.21	4.8

Not Ready Breakdown

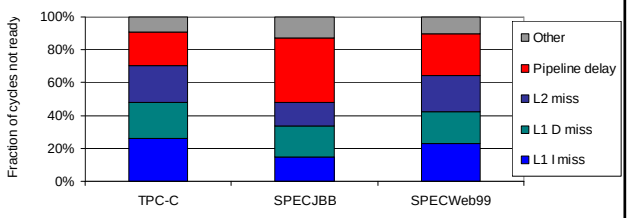

- TPC-C - store buffer full is largest contributor
- SPEC-JBB - atomic instructions are largest contributor
- SPECWeb99 - both factors contribute

Performance: Benchmarks + Sun Marketing

Benchmark\Architecture	Sun Fire T2000	IBM p5-550 with 2 dual-core Power5 chips	Dell PowerEdge
SPECjbb2005 (Java server software) business operations/ sec	63,378	61,789	24,208 (SC1425 with dual single-core Xeon)
SPECweb2005 (Web server performance)	14,001	7,881	4,850 (2850 with two dual-core Xeon processors)
NotesBench (Lotus Notes performance)	16,061	14,740	

SPECjappServer 2004 Dual Node	Sun Fire T2000	HP rx4640
Space (RU)	2	4
Watts	320	1,303
Performance (SPECjapp JOPs)	615	471
Performance / Watt	1,922	0.361
SWaP	0.96	0.09

Space, Watts, and Performance

Performance: Benchmarks + Sun Marketing

Benchmark\Architecture	Sun Fire T2000	IBM p5-550 with 2 dual-core Power5 chips	Dell PowerEdge
SPECjbb2005 (Java server software) business operations/ sec	63,378	61,789	24,208 (SC1425 with dual single-core Xeon)
SPECweb2005 (Web server performance)	14,001	7,881	4,850 (2850 with two dual-core Xeon processors)
NotesBench (Lotus Notes performance)	16,061	14,740	

SPECjappServer 2004 Dual Node	Sun Fire T2000	HP rx4640
Space (RU)	2	4
Watts	320	1,303
Performance (SPECjapp JOPs)	615	471
Performance / Watt	1,922	0.361
SWaP	0.96	0.09

Note the paradigm shift

Space, Watts, and Performance

Microprocessor Comparison

Processor	SUN T1	Opteron	Pentium D	IBM Power 5
Cores	8	2	2	2
Instruction issues / clock / core	1	3	3	4
Peak instr. issues / chip	8	6	6	8
Multithreading	Fine-grained	No	SMT	SMT
L1 I/D in KB per core	16/8	64/64	12K uops/16	64/32
L2 per core/shared	3 MB shared	1MB / core	1MB/ core	1.9 MB shared
Clock rate (GHz)	1.2	2.4	3.2	1.9
Transistor count (M)	300	233	230	276
Die size (mm^2)	379	199	206	389
Power (W)	79	110	130	125

Performance Relative to Pentium D

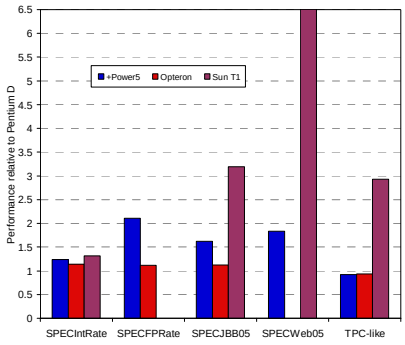

Performance/mm^2, Performance/Watt

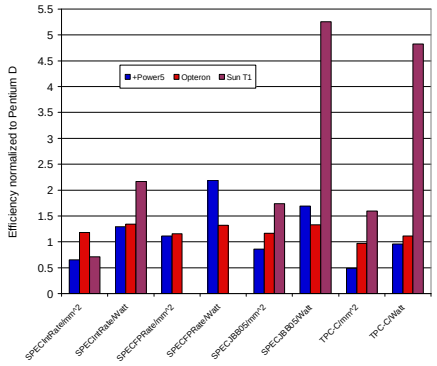

Niagara 2

- Improve performance by increasing threads supported per chip from 32 to 64
  - 8 cores * 8 threads per core
- Floating-point unit for each core, not for each chip
- Hardware support for encryption standards EAS, 3DES, and elliptical-curve cryptography
- Niagara 2 will add a number of 8x PCI Express interfaces directly into the chip in addition to integrated 10Gigabit Ethernet XAU interfaces and Gigabit Ethernet ports.
- Integrated memory controllers will shift support from DDR2 to FB-DIMMs and double the maximum amount of system memory.

Kevin Krewell
"Sun's Niagara Begins CMT Flood - The Sun UltraSPARC T1 Processor Released"
Microprocessor Report, January 3, 2006

Sun Niagara 2 at a Glance

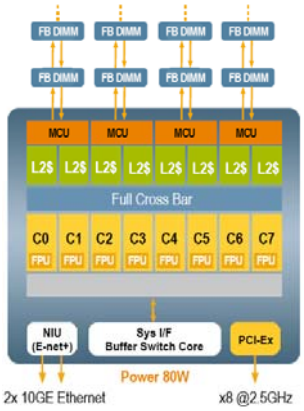

- 8 cores x 8 threads = 64 threads
- Dual single issue pipelines
- 1 FPU per core
- 4MB L2, 8-banks, 16-way S.A
- 4 x dual-channel FBDIMM ports (60+ GB/s)
- > 2x Niagara 1 throughput and throughput/watt
- 1.4 x Niagara 1 int
- > 10x Niagara 1 FP
- Available H2 2007

Sparc T3

Third Generation Niagara Processer

SPARC T3 Architecture

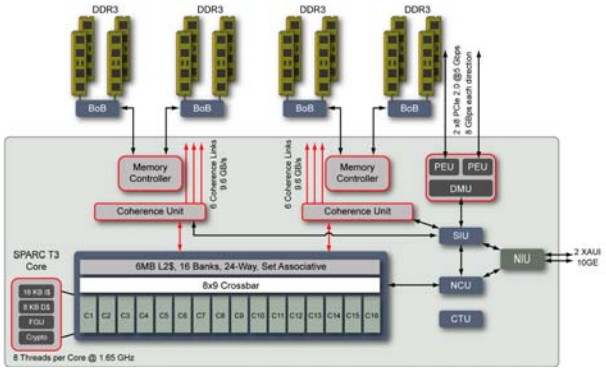

SPARC T3 Pipelines

Fetch	Cache	Pick	Decode	Execute	Mem	Bypass	W

Eight-Stage Integer Pipeline

Fetch	Cache	Pick	Decode	Execute	Fx1	Fx2	Fx3	Fx4	Fx5	Fx6	FW

Twelve-Stage Floating-Point Pipeline

SPARC T3 Thread Management

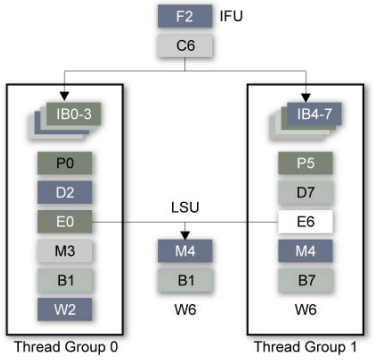

SPARC T3 Multi-level Threading

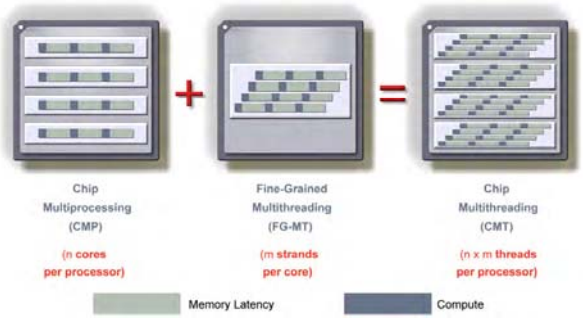

SPARC T3 Threading

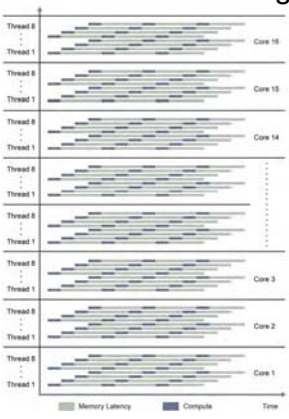

Performance Comparison—SPECweb2005

Competitive Comparisons Table

	SPARC T3-2	Fujitsu PRIMERGY TX300 S6	HP DL585 G7	IBM BladeCenter HS22
Processor	SPARC T3	Intel Xeon X5680	AMD Opteron 6128 HE	Intel Xeon X5570
clock	1.66 GHz	3.33 GHZ	2 GHZ	2.94 GHz
cores	16	6	8	4
threads/core	8	2	1	2
TDP	139 W	130 W	85 W	95 W
# of Procs.	2	2	4	2
# of Threads	256	24	32	16
Op. Sys.	Oracle Solaris 10	Red Hat Linux	Red Hat Linux	Red Hat Linux
Web Server	Oracle iPlanet 7	Rock Web Server v1.4.8	Rock Web Server v1.4.9	Rock Web Server v1.4.9
SPECweb2005 Performance (Composite)	**113,857**	**104,422**	**105,586**	**75,155**