



Gemini:

A Power-efficient Chip Multi-Threaded (CMT)
UltraSPARC® Processor

Sanjiv Kapil

Gemini Architect

Sun Microsystems, Inc.





Design Goals


- **Designed for compute-dense, transaction oriented systems (webservers, appservers)**
 - **Small form factors (1U, 2U, Blades)**
 - **Low power, low cost**
 - **Large physical memory (64-bit addressing)**
 - **Enterprise-class RAS**
- **Focus on throughput performance rather than single thread performance**
- **Support for 1- to 4-way SMP with simultaneous execution of 2-8 threads**
- **Maintains binary compatibility with SPARC® V9 code**

UltraSPARC Processor Roadmap

Data Facing



Network Facing



Features

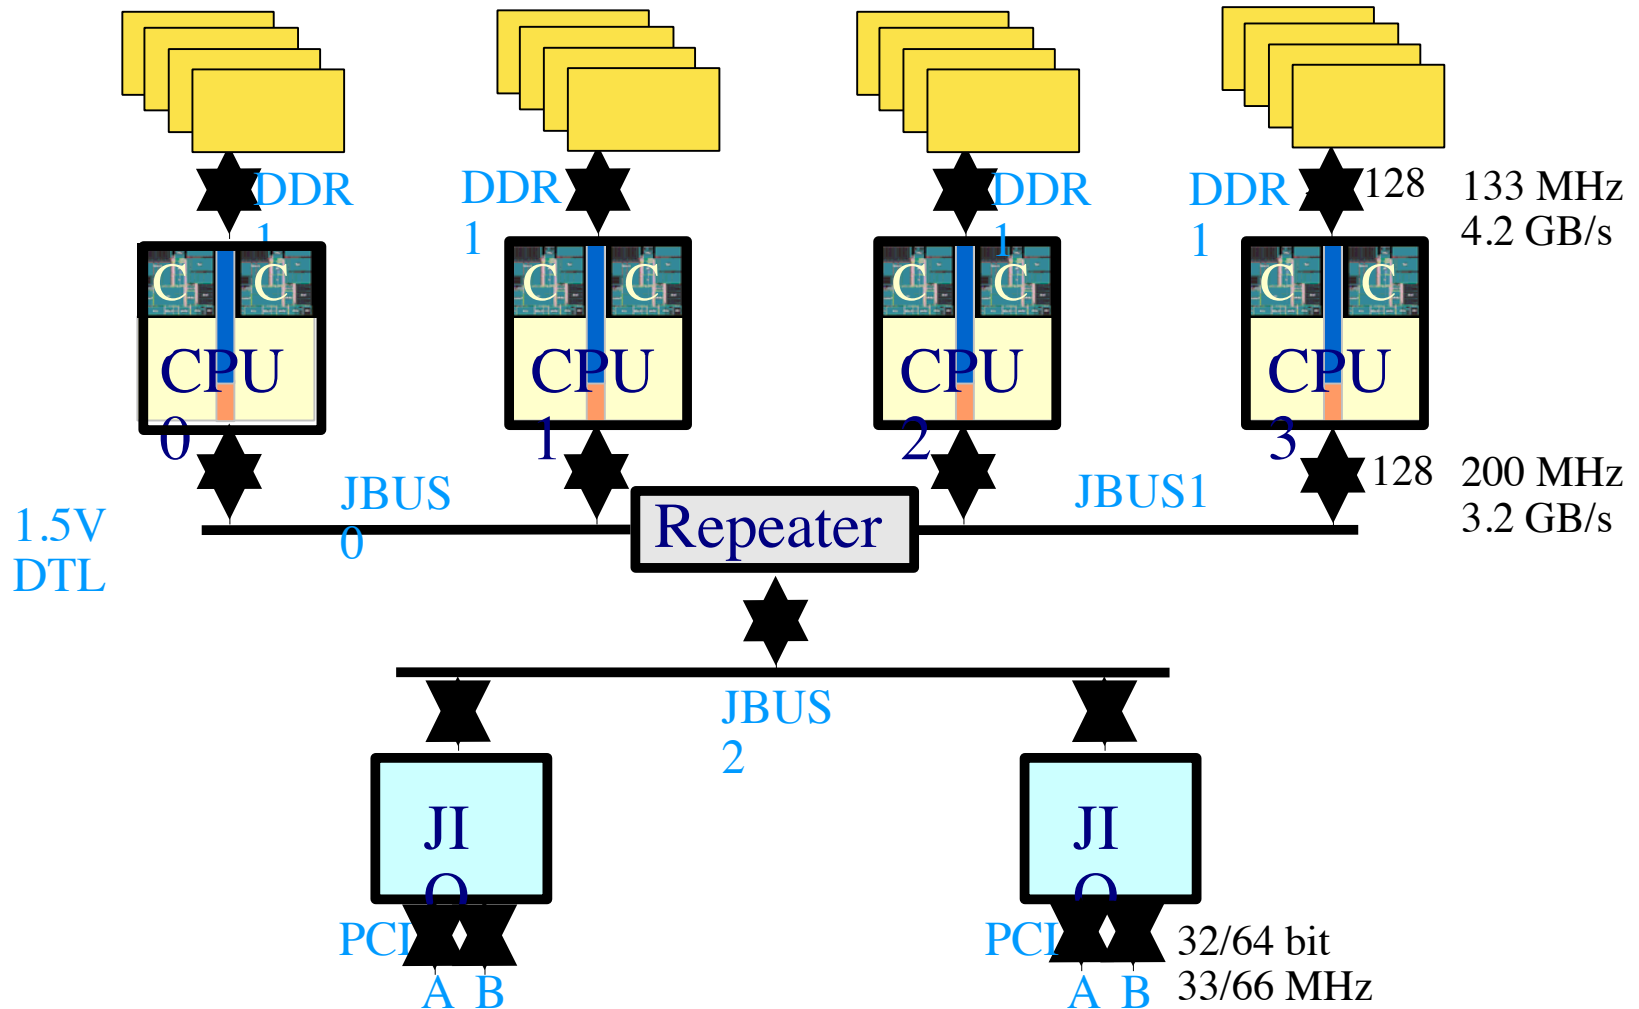
- *Dual 4-issue superscalar cores*
 - *Based on 64-bit SPARC V9 processor architecture*
 - *2 integer ops, 2 FPU/graphics (VIS 1.0) ops/cycle*
- *Dual integrated 0.5MB L2 caches (one per core)*
 - *4-way set associative*
 - *Way-locking per core*
- *Pin-compatible with UltraSPARC IIIi processor*
 - *JBus system bus interface*
- *Integrated memory interface to standard DRAM*
 - *Shared 128-bit DDR1 controller*
 - *Up to 16 GBs total using 2-4 DIMMS*

Features (con.)

- **Power management modes (E*Star)**
 - Transition to either $F_{cpu}/2$ or $F_{cpu}/32$
- **Robust RAS support**
 - Parity checking on L1 caches, JBus
 - ECC on L2 caches, DRAM
- **CMT features**
 - Software-controlled core disable
 - Dynamic cache-aware core parking
 - Adheres to UltraSPARC CMT Architectural Specification


Gemini-based System

4-Way SMP/8-thread System




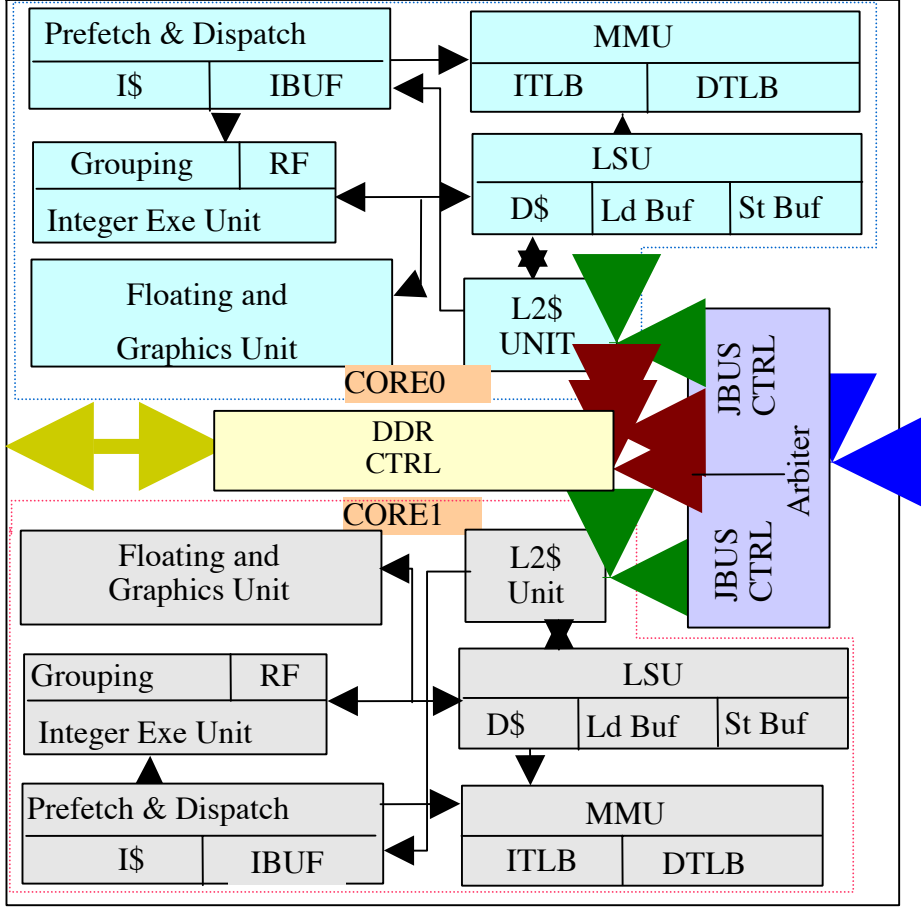
Technology

- **Dual-core CMT design**
- **Transistor count: 80 million**
- **Die size: 206 sq. mm.**
- **Package: 959 pin ceramic uPGA**
- **Maximum power consumption: 32 Watts @ 1.2 GHz, 1.3V**
- **Technology:**
 - **TI's advanced 130 nm CMOS process, 300 mm wafers**
 - **53 nm L_{eff}**
 - **7LM Cu, low-k dielectric (OSG)**
 - **dual V_t**




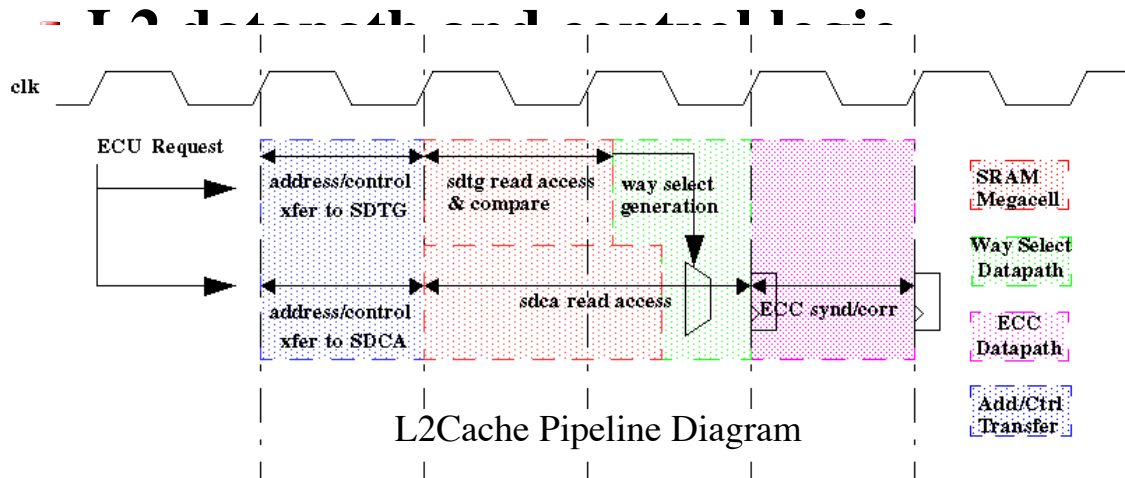
Block Diagram

Dual Core 64Bit Superscalar Processor



L2 Caches

- Unified 512KB L2 cache per core: 4-way set associative, 64B line size, Pseudo-random replacement
- 2-2 mode access with 4-cycle latency
- 8 ECC bits for 64b data with physical bit scrambling, parity-protected tags



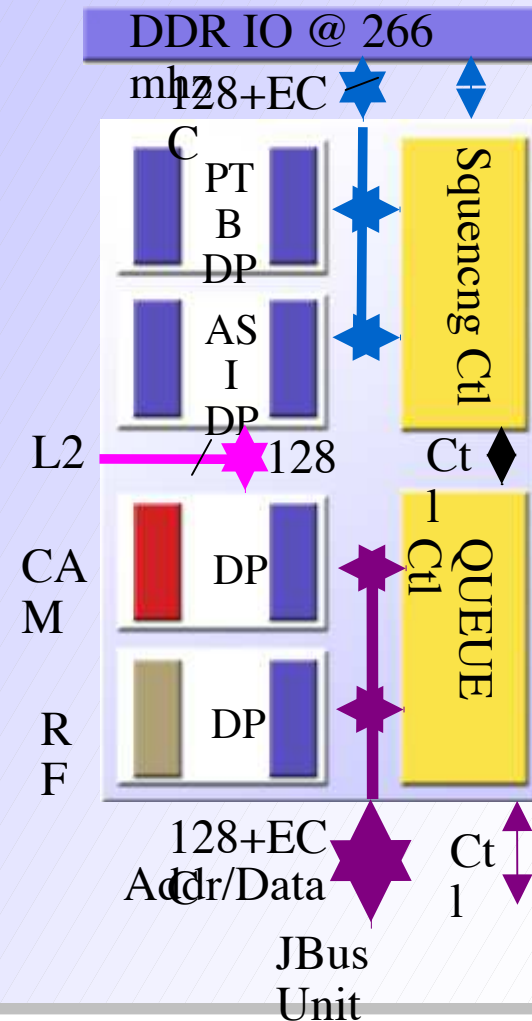
Memory Controller

Specification

S

- Shared controller for two cores
- DDR1 compliant
- Max 4 DIMMs (single-/double-sided)
- Total memory 256MB to 16GB
- SSTL Input/Output
- Memory I/F: 128-bit data + 9 ECC check bits
- E*Star support (1/2 or 1/32 of F_{cpu})
- PTB to track 16 open pages
- 4.2 GB/s peak B/W
- Measured local access latency: 96ns


DDR Unit




Memory Subsystem Details

Specifications	Main Features
<p>Caches: I: 16KB 2-way associative PIPT with 32B line size D: 16KB direct VIPT with 32B line size L2: 512 KB, 4-way set associative 64B line size Pseudo-random replacement PIPT</p>	<p>I: Parity protection on tags/data s/w recoverable D: Parity protection on data and tag L2: Parity protection on tags, ECC on data 2-2 mode access 3-ways locking support Flush to memory support</p>
<p>Memory: DDR1: 2.5 v (2.625v tolerant) SSTL 128-bit data + 9-bit ECC Peak memory B/W: 4.2GB/@133Mhz 128/256/512/1G bit 2-4 Dimm support 64-bit virtual address space</p>	<p>Shared for dual cores PTB tracks 16 open pages E*Star modes capable Separate PLL for DDR/E*Star support 4-Way SMP can address 64GB space</p>
<p>CMP: Core disable/parking/error reporting JTAG/RAMTEST access to CSR's</p>	<p>Non-core/Core-specific error reporting Full 16 GB addressing by each core Cache-aware parking (No instr. Execution)</p>


Area Distribution




Core Area = 28.6 mm²




CPU Area = 206 mm²




Power Distribution




Core Power = 6.7 W



CPU Power = 29.3 W



UltraSPARC Design Evolution



16,508 trans/mm²
315 mm²

UltraSPARC I – 1995

5.2M transistors
500 nm CMOS
4 LM Al
3.3V, 30W @ 167 MHz
SPEC2K int/rate: 76/0.9*

36,242 trans/mm²
149 mm²

UltraSPARC II – 1997

5.4M transistors
350 nm CMOS
5 LM Al
2.5V, 25W @ 250 MHz
SPEC2K int/rate: 120/1.4*

388,350 trans/mm²
206 mm²

Gemini – 2003

80M transistors
130 nm CMOS
7 LM Cu
1.3V, 32W @ 1200 MHz
SPEC2K int/rate: 511/11
4x perf, 8x throughput

*calculated from measured SPEC95 scores

Performance

- **Peak 3.2 GB/s system I/O bandwidth**
- **Peak 4.26 GB/s memory bandwidth**
- **Average memory latency for local DRAM 96ns**

CPU's	Int_Rate	/Watt
1	11	0.34
2	16.2	0.25

CPU's	FP_Rate	/Watt
1	12.2	0.38
2	18.0	0.28

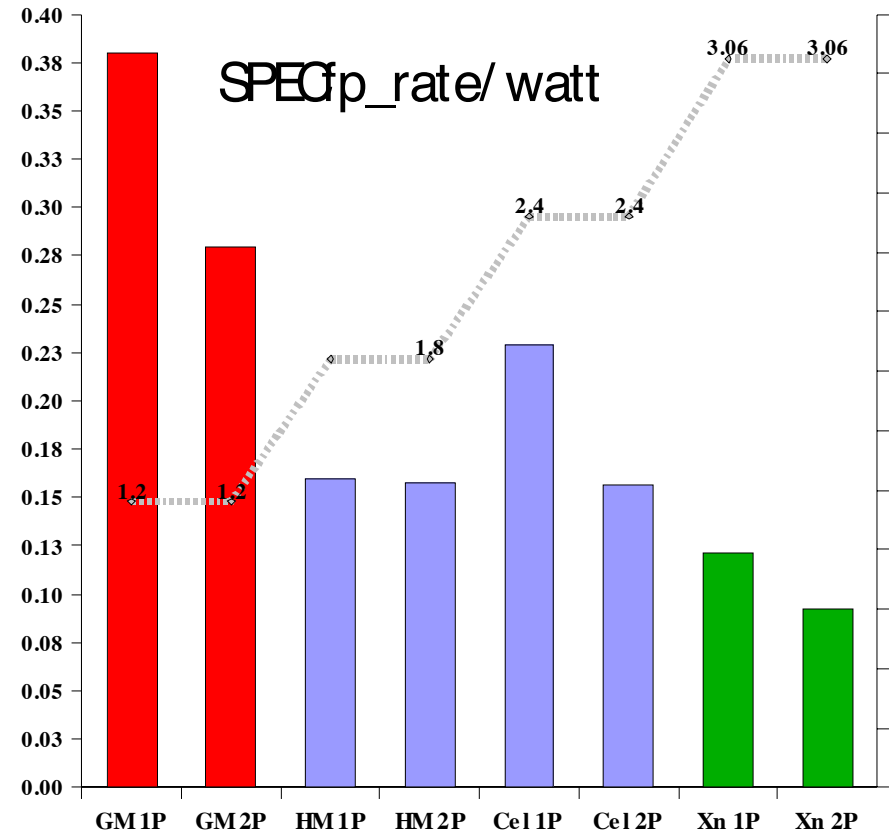

CPU's	SPECjbb2k	/Watt
1	18700	584
2	34500	539

CPU's	SPECweb99	/Watt
1	2100	66
2	3400	53

All scores @ 1.2 GHz with 16 GB memory,
2 active cores/processor


All numbers preliminary based on measured data

Benchmark: SPECint/fp_rate



Sources: www.spec.org; AMD Opteron Processor Data Sheet – Apr. 2003; Intel Celeron Processor Data Sheet (2-2.6 GHz) – Jun. 2003; Intel Xeon Processor Data Sheet (2-3.06GHz) – Jul. 2003

Benchmark: SPECweb, SPECjbb



Sources: www.spec.org; AMD Opteron Processor Data Sheet – Apr. 2003; Intel Celeron Processor Data Sheet (2-2.6 GHz) – Jun. 2003; Intel Xeon Processor Data Sheet (2-3.06GHz) – Jul. 2003

Summary

- **First generation CMT design for high-density, low-cost, network-facing systems**
- **Optimized for throughput rather than single-thread performance**
- **Maintains binary compatibility with existing SPARC code base**
- **Highly leveraged design methodology**

Q & A