The UltraSPARC® IIIi Processor Architecture Overview

Technical Whitepaper **Version 1.2, April 2004**

This document provides a brief overview of the important new technology of the UltraSPARC® IIIi processors, specifically as it serves to maintain compatibility, accelerate real performance, and improve RAS (Reliability, Availability, Serviceability) in SPARC-based systems designed for high-end desktop workstations and workgroup servers.

4150 Network Circle
Santa Clara, CA 95054 USA
http://www.sun.com

Table of Contents

CHAPTER 1

Overview

The UltraSPARC® IIIi processor is a high performance, highly-integrated 4-issue superscalar processor implementation of the 64-bit SPARC®-V9 RISC architecture, supporting a 64-bit virtual address space. The core instruction set has been extended to include specialized instructions that provide the most common operations related to graphics processing, data compression, signal processing, 3D visualization, networking, telecommunications, and other applications.

1.1 Key Characteristics

Processor

- SPARC V9 Architecture with the VIS Instruction Set; SPARC binary code compatible
- Based on the UltraSPARC III processor
- 64-bit virtual address space
- 43-bit physical address space
- 1.28 GHz
- Multiprocessor system capability - 1 to 4 processors
- On-chip L2 Cache & Translation Lookaside Buffers
- Fast, wide 64-bit data paths
- 64-bit ALUs
- Steers up to 4 instructions per cycle into 6 execution pipes
- 14 stage, non-stalling pipeline
- Low cost, low power design (56W@1.28GHz)
- 32KB, 4-way set associative L1 instruction cache
- 64KB, 4-way set associative L1 data cache
- 2KB prefetch cache
- 2KB write cache
- 1MB 4-way set associative, unified instruction/data on-chip L2 cache

Technology

- 130nm, 7-layer Cu metal, CMOS process
- 1.4V core, 3W I/O power
- 183mm^2 die size
- 56W power dissipation at 1.28GHz
- 959 pin ceramic Pin Grid Array (cPGA) package
- 469 signal pins (excludes spares and not connected signals)
- 87.5M transistors (of which 63M used for L2 cache)

Integration

- JBUS interconnect interface
  - 200MHz operation
  - 16-byte shared address/data bus
  - 3.2GB/s peak bandwidth
- DDR Memory interface
  - 4.256GB/s peak bandwidth
- Clock Control Unit: enables low power systems
  - 1/2 and 1/32 reduced power dissipation

1.2 Key Applications

The density and versatility of the UltraSPARC IIIi processor makes it a very good choice for many applications from desktop workstations and small enterprise servers to blade servers. The primary focus is enterprise applications that require scalable performance and can take advantage of multi-processing, larger on-chip cache, larger memory, and increased I/O bandwidth.

The configuration of the UltraSPARC IIIi processor makes it suitable for both throughput-oriented and transactional-oriented workloads. The 4-way system is ideal for applications such as:

- On-line transaction processing (OLTP)
- Media encoding and decoding
- Data Warehousing/Mining
- E-Commerce
- Messaging (E-mail)
- Graphics and image applications
- Networking
- Enterprise Resource Planning (ERP)

CHAPTER 2

Design Goals

The UltraSPARC IIIi processor shares the same four basic design goals as the other members of the UltraSPARC processor family including: compatibility, scalability, performance and RAS.

2.1 Compatibility

Binary compatibility is the primary design goal of any new SPARC processor. Typically, legacy applications will benefit in terms of performance without having to recompile. Furthermore, this benefit applies to all applications.

2.2 Scalability

SPARC processors are designed to scale up gracefully in performance not just from generation to generation, or implementation to implementation, but also from single processor configuration to multiprocessor configuration. Scalability also requires architecting the on-chip memory system and the bus interface to handle multiprocessor systems to be built using the UltraSPARC IIIi processor. Initial configuration of the system bus, JBUS, can support up to four UltraSPARC IIIi processors.

2.3 Performance

Another key design goal of any new UltraSPARC processor design is performance. For the UltraSPARC IIIi processor, the performance target shifts from high-end desktop systems and powerful multi-way servers to mid-range desktop systems and cost-effective, small-way servers. The UltraSPARC IIIi processor achieves the highest possible real performance when used by customers to do actual work in normal business environments, running various applications.

The UltraSPARC IIIi processor offers large L1 instruction and data caches, large on-chip L2 cache, an on-chip memory controller that supports DDR memory up to 16GB. To support multiprocessor capabilities, the system has a new cache-coherent, Symmetric Multi-Processor SMP[1] system bus interface that can achieve 3.2GB/s bus bandwidth.

2.4 RAS Features

Error Detection and Correction (EDC) logic is used extensively throughout the UltraSPARC IIIi processor data and address paths to provide high levels of data integrity. In both the L2 cache and main memory, all single-bit errors are corrected and double-bit errors are detected.

In addition, the UltraSPARC IIIi processor provides parity protection on the physical tag array, snoop tag (sTag) array and data array of the L1 I-cache and L1 D-cache. If a parity error is detected, recovery is possible by invalidating the affected cache line and reloading it from a lower level of memory when the line is needed. The L2 cache (including tags and data) and main memory are all protected by ECC.

In general, the UltraSPARC IIIi processor is able to eliminate more errors than any previous generation of UltraSPARC processors. In small-scale MP systems, each UltraSPARC IIIi processor controls its own local memory thus minimizing the threat of a data error propagating to other processors in the system.

CHAPTER 3

Key Features

The design of the UltraSPARC IIIi processor is based on the third-generation SPARC V9 64-bit architecture. The processor has a 14-stage nonstalling pipeline that allows the concurrent execution of four instructions per cycle in 6 execution units: 2 integer for arithmetic, logic and shift operations, 2 floating point units for addition and multiplication, 1 unit for memory load/store, and 1 unit for branch instructions.

The translation look-aside buffer supports 8KB, 64KB, 512KB, and 4MB pages. On-chip L1 caches include a 64KB data cache, a 32KB instruction cache, a 2KB data prefetch cache, and a 2KB write-cache. All caches are 4-way set associative. The instruction and data caches have parity protection.

The new features in this microprocessor include an optimized system bus interface tailored to the needs of 1 to 4 processor high-end desktop and workgroup servers, with focus on higher system integration and cost reduction. This new system bus interface, called JBUS, is a 200MHz cache coherent SMP bus interface. An on-chip L2 cache, error-correction-code (ECC) protection on data, improves performance and system level integration. The microprocessor also includes a DDR memory interface. The main functional block diagram is shown in Figure 3-1.

3.1 Cache, Memory Controller and System Bus Interface

3.1.1 L2 Cache

The L2 cache is 1MB with 64-byte lines. Additionally, the L2 cache is 4-way set associative and physical indexed, physical tagged. It is a write-back cache, supporting modified-owned-exclusive-shared-invalid (MOESI) cache coherency protocol. The cache has parity protection on tag and ECC protection on data allowing correction of any single-bit error and detection of any double-bit errors.

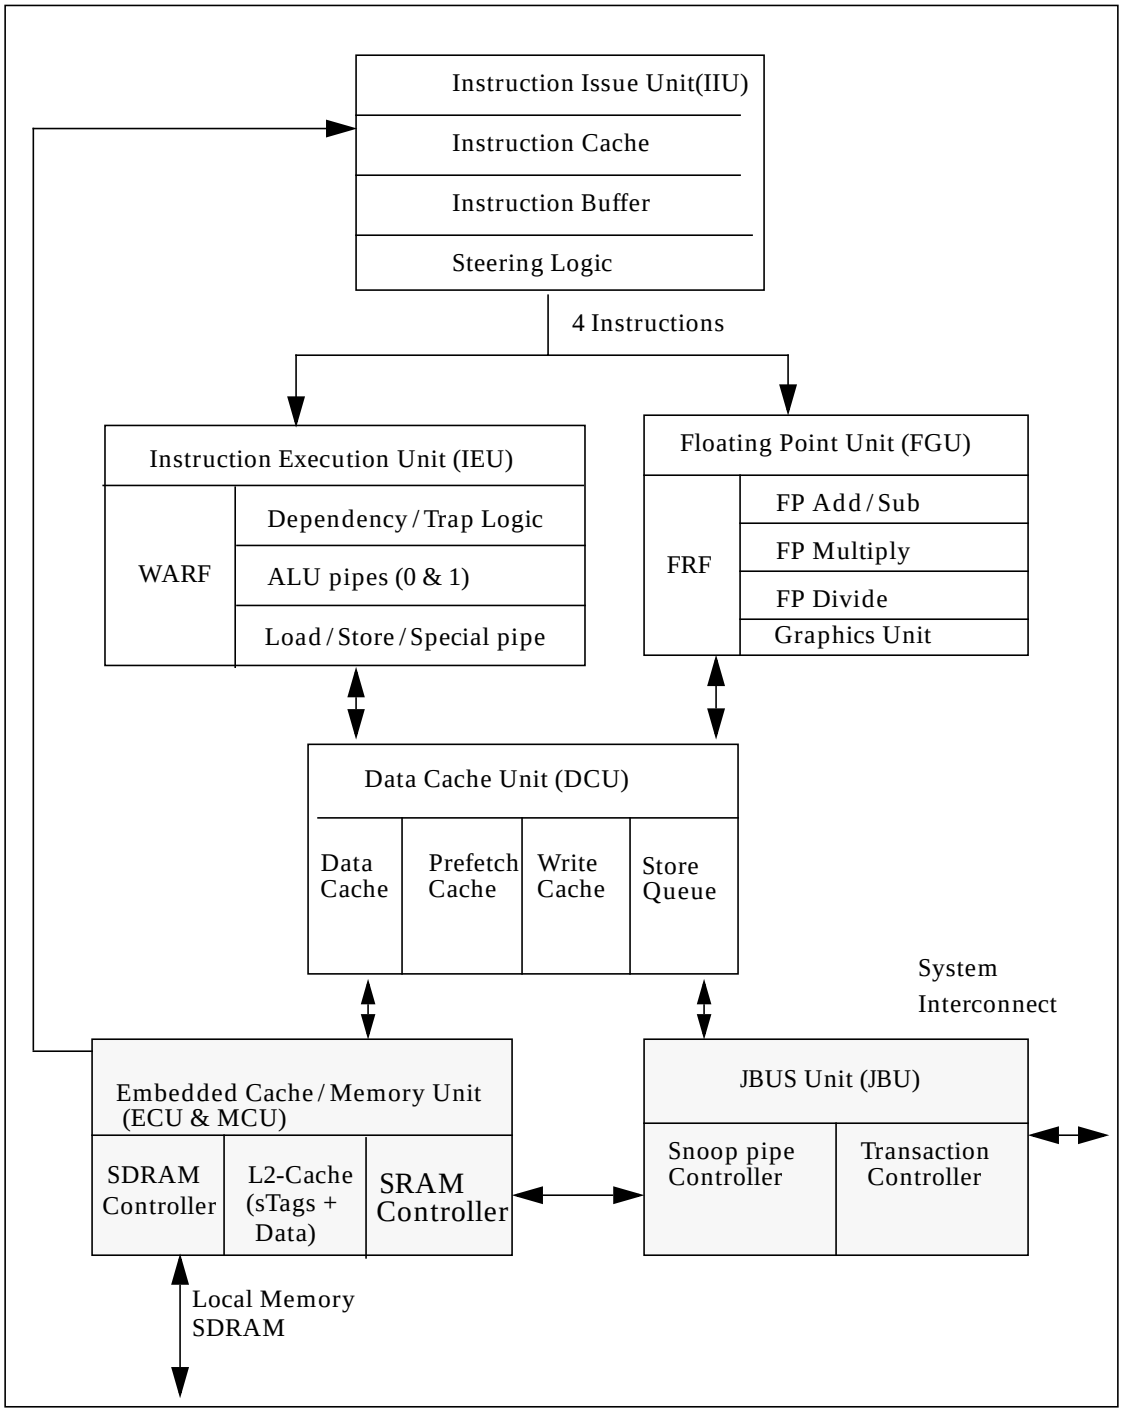

Figure 3-1 The UltraSPARC IIIi Processor Block Diagram

3.1.2 Memory Controller

Another architectural feature of the UltraSPARC IIIi processor is its on-chip memory controller. The Memory Control Unit controls the operation of the external memory. Moving the memory controller on-chip reduces memory latency, and scales memory bandwidth with the number of processors. The controller can support up to 4 banks of DDR memory, totaling 16GB. The controller has a transfer rate of 4.256GB/s. The UltraSPARC IIIi processor marks the world's first integration of elements of asynchronous logic technology developed at Sun Laboratories. The memory controller includes asynchronous logic-based first-in, first-out (FIFO) circuits in the memory controller I/O section to absorb clock skew variations inherent in large transistor designs.

3.1.3 System Bus Interface

The new system bus interface, called JBUS, is a cache-coherent packet-switched bus that supports multiple agents or ports, with a peak bandwidth of 3.2GB/s. The bus supports distributed round-robin arbitration with snooping that is fully decoupled from memory access.

3.2 Physical Design Implementation

Like many previous UltraSPARC processors, the UltraSPARC IIIi processor is fabricated by Texas Instruments (TI). The following table summarizes the physical implementation of the UltraSPARC IIIi processor.

Table 3-1 UltraSPARC IIIi Implementation

Parameter	**UltraSPARC IIIi**
Process	130nm CMOS
Frequency (GHz)	1.28
Metal Layers	7 copper
Transistors (millions)	87.5
Core Voltage (volts)	1.4
Peak Pwr. Dissipation (watts)	56@1.28GHz
Die Area (mm^2)	183
Pins	959
Package	ceramic PGA
L2 cache	1MB on-chip

CHAPTER 4

Multiprocessor Systems

The UltraSPARC IIIi processor is designed to create a 1-4 way multiprocessing systems. Any system constructed with the UltraSPARC IIIi processor must contain a JBUS and a JBUS ASIC. The JBUS is a system bus that supports small-scale 1-4 way multiprocessor system by implementing a fast 'glueless logic' interconnect between processors. The JBUS ASIC is a companion to the UltraSPARC IIIi processor. The JBUS ASIC provides the interface to external I/O and graphics for the UltraSPARC IIIi processor.

4.1 JBUS Interconnect

JBUS delivers the high bandwidth and low latency needed for networking, communications and other embedded applications. Built for the bandwidth-hungry Internet infrastructure, JBUS targets 64-bit multiprocessor servers. It features a 128-bit packet-switched split-transaction request and data bus. It is a flexible, extensible and easily implemented bus structure. Processors can attach to a coherent shared bus without any glue logic. JBUS is capable of providing scalable 1-4 way SMP configurations with a selection of data bus widths, power, space and cost options.

4.2 JBUS ASIC

The JBUS ASIC is a companion to the UltraSPARC IIIi processor. The JBUS ASIC and the processor(s) communicate through a high speed split-transaction 16-byte shared address/data JBUS. The primary goal of the JBUS ASIC is to provide the interface to external I/O and graphics for a SMP 1-4 way system. It also offloads the I/O traffic from the JBUS to provide faster inter-processor communication.

The JBUS ASIC has five major functions: A JBUS Cluster, a PCI Cluster, a Graphics Cluster, an I/O cache and an I/O Memory Management Unit (IOMMU). The graphic interface may be configured to operate in one of two possible modes: PCI or UPA64s (UltraSPARC Port Architecture). The selection of which mode is active will be fixed via motherboard design, and thus not changeable during system operation. The following section gives a more detail description on each of the JBUS ASIC functions.

4.3 Implementation Examples

4.3.1 Uni-Processor System

Figure 4-1 illustrates a single-processor, UltraSPARC IIIi-based system with one JBUS ASIC. This is the minimal configuration. Any system constructed with JBUS must contain at least one JBUS ASIC. Note that the figure shows the UltraSPARC IIIi processor minimum memory configuration of only two Dual In-Line Memory Modules (DIMMs) on the 16-byte wide Double Data Rate (DDR). In this one processor configuration, the low cost PCI is used as the I/O bus interface

4.3.2 Dual Processor Systems

Figure 4-2 shows an example of a dual UltraSPARC IIIi processor system, with two JBUS ASICs for expanded I/O. The JBUS supports the communication, memory sharing, and the cache coherency between processors. The link to the I/O is via the JBUS ASIC. This example shows the maximum memory of four DIMM chips per processor.

.

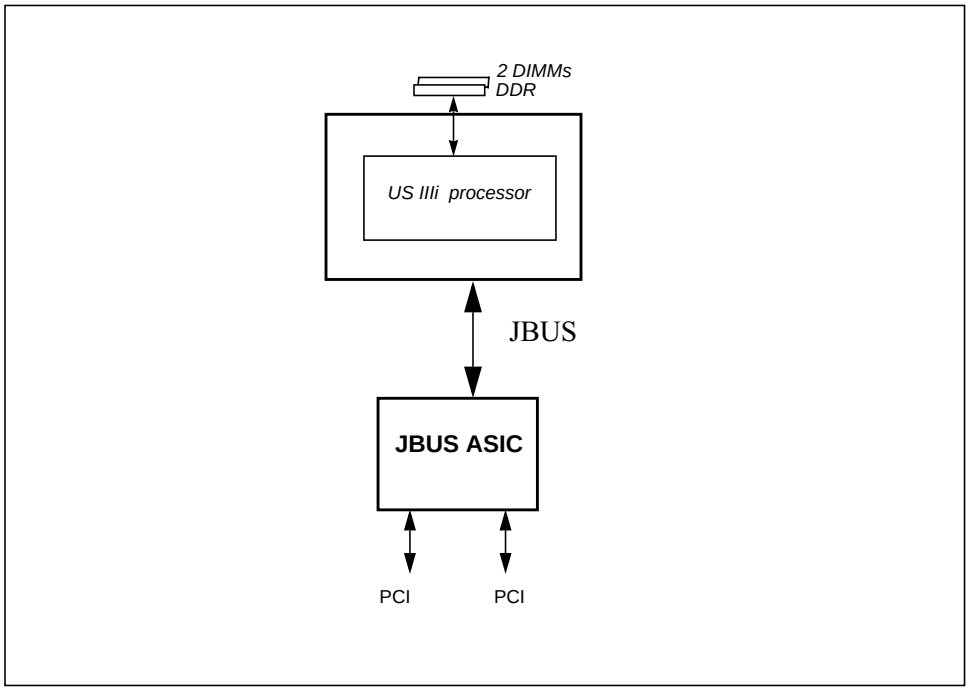

Figure 4-1 One-way -UltraSPARC IIIi Processor and JBUS Interconnect

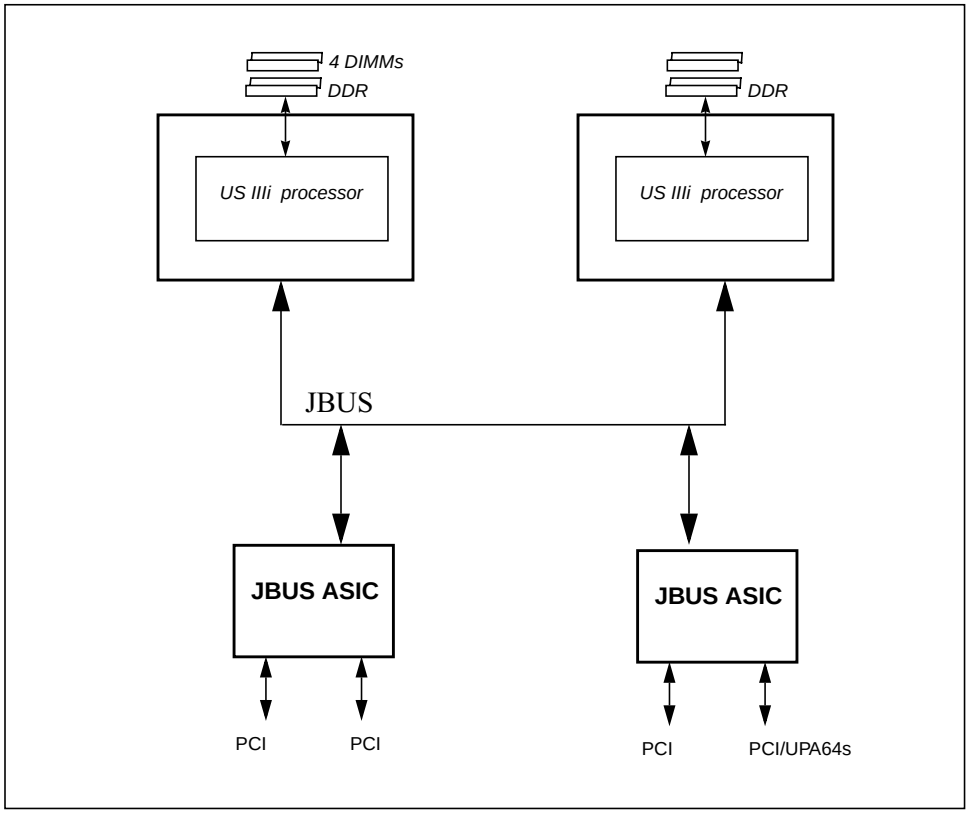

Figure 4-2 Two-way UltraSPARC IIIi Processors and JBUS Interconnect

4.3.3 4-Processor System

Figure 4-3 shows an example of a JBUS-based 4-way UltraSPARC IIIi processor system. Up to 4 UltraSPARC IIIi processors can be supported by the JBUS ASIC. But in the configuration shown, the repeater chip is inserted to enable JBUS to run at its peak frequency. Two JBUS ASICs are used to provide more I/O.

As shown in the figure, the JBUS ASIC has a multi-purpose interface that supports both UPA64S (UltraSPARC Port Architecture) and PCI buses. This allows either use of Sun's high-performance graphics accelerators or external PCI graphics devices on the same motherboard design.

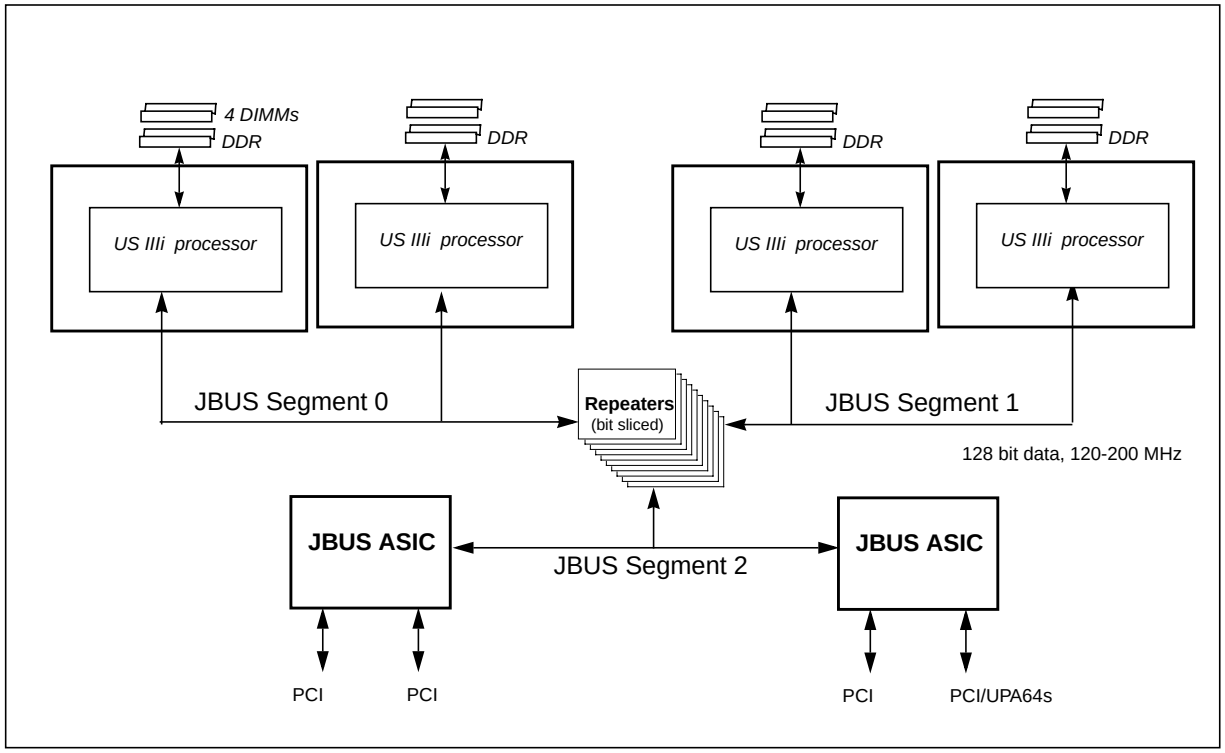

Figure 4-3 Four-way UltraSPARC IIIi Processors and JBUS Interconnect

CHAPTER 5

Summary

The UltraSPARC IIIi processor is the new generation of 64-bit SPARC processor to be used in a wide range of Sun systems. The architecture enables multiprocessor capabilities with up to 4 processors to be easily built and achieve excellent performance. Specifically, the UltraSPARC IIIi processor has been designed to provide software compatibility with all previous generations of SPARC design, both 32-bit and 64-bit.

The features of the UltraSPARC IIIi processor that contributes to its leading-edge performance include:

- Large on-chip L2 cache delivers high bandwidth for performance demanding applications such as data warehousing, on-line transaction processing, or technical computing.
- 4-issue superscalar processor.
- On-chip memory controller supports 16GB of DDR memory, and provides 4.256GB/s per processor memory bandwidth.
- The world's first integration of asynchronous logic elements developed at Sun Laboratories. The processor's memory controller includes asynchronous logic circuits in the memory controller I/O section to absorb clock skew variations.
- Synchronous, packet-switched, cache-coherent, JBUS system interface enabling multiprocessors to communicate via a shared 3.2GB/s bus.
- Specialized instructions to accelerate specific processing tasks for multimedia, scientific, image processing, Java, and network applications.

The UltraSPARC IIIi processor is designed to be versatile, powerful, and reliable to meet the needs of many different types of applications. The UltraSPARC IIIi processor delivers impressive computing throughput in a highly integrated package, with efficient power usage to make it applicable to a wide-range of compact, low-cost systems.

References

A. Background

1. D. E. Culler, J.P. Singh and A. Gupta, Parallel Computer Architecture, Morgan Kaumann Publisher, Inc., San Franciso, California, 1999.

2. J. Duato, S. Yalmanchili and L. Ni, Interconnection Networks: An Engineering Approach, IEEE Computer Society Press, 1997.

3. K. Hwang and Z. Xu, Scalable Parallel Computing, McGraw-Hill, 1998.

4. K. Hwang, Advanced Computer Architecture: Parallelism, Scalability, programmability, McGraw-Hill, 1993.

5. D. E. Lenoski and W.D. Weber, Scalable Shared-Memory Multiprocessing, Morgan Kaumann Publisher, Inc., San Franciso, California, 1995.

6. D. Patterson and J. L. Hennessy, Computer Organization and Design: The Hardware/Software Interface, 2nd Edition, Morgan Kaumann Publisher, Inc., San Franciso, California, 1997.

B. Architecture of JBUS

7. R. N. Raman, “JBUS: A Bus Architecture for Embedded On-chip Multiprocessing”, Embedded Microprocessor Forum, 2001.

C. UltraSPARC IIIi Processor

8. K. Normoyle, “A Dash of Jalapeno: Introducing the UltraSPARC IIIi Microprocessor”, Microprocessor Forum, 2001.

9. G. Konstadinids et al., “Implementation of a Third-Generation 1.1-GHz 64-bit Microprocessor,” IEEE Journal of Solid-State Circuits, Nov. 2002, pp. 1461-1469.