

A 1.3GHz Fifth Generation SPARC64 Microprocessor

Hisashige Ando, Yuuji Yoshida, Aiichiro Inoue, Itsumi Sugiyama, Takeo Asakawa,
 Kuniki Morita, Toshiyuki Muta, Tsuyoshi Motokurumada, Seishi Okada, Hideo Yamashita,
 Yoshihiko Satsukawa, Akihiko Konmoto, Ryouichi Yamashita, Hiroyuki Sugiyama
 Fujitsu Ltd. Kawasaki, Japan 211-8588
 hando@jp.fujitsu.com, {yoshida, inoue}@cs.fujitsu.co.jp

ABSTRACT

A 5th generation SPARC64 processor is fabricated in 130nm SOI CMOS process with 8 layers of Cu metallization. It runs at 1.3GHz with 34.7W power dissipation in the laboratory. The chip contains over 190M transistors with 19M in logic circuits. The chip size is 18.14mm x 15.99mm. The error detection and recovery mechanism is implemented for execution units and data path logic circuits in addition to on-chip arrays to detect and recover from data logic error. This processor is developed by using mostly in-house CAD tools.

Categories and Subject Descriptors

B7.1 [INTEGRATED CIRCUITS] Types and Design Styles
 Subjects: Microprocessors and microcomputers, B8.1 [Performance and Reliability]: Reliability, Testing and Fault-Tolerance.

General Terms

Design, Reliability.

Keywords

Microprocessor, SPARC, microarchitecture, reliability, clock distribution, unix server.

1. INTRODUCTION

This 5th generation SPARC64 processor is designed for high end unix servers. Business mission critical servers must operate 24 hours x 7 days reliably. A high data integrity is especially important since an undetected data error is the worst thing to happen that cast doubt on all outputs from the server. Since high end servers use many processor chips, the MTBF requirement for each processor is much more stringent than the one for single processor workstations. That is why this processor is designed to detect errors in execution units and data paths and also to recover from those detected errors as much as practically possible. These high reliability requirements were not fully met with the previous processors designed for unix servers.

The improvement in performance without consuming too much power is the important goal of this processor development. The target was to achieve twice the performance with the equal or less power compare to the then current 4th generation SPARC64 processor which runs at 563MHz and consumes about 50W. To achieve

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee

DAC 2003, June 2-6, 2003, Anaheim, California, USA.
 Copyright 2003 ACM 1-58113-688-9/03/0006...\$5.00.

this goal, a new 130nm partially depleted SOI CMOS process is selected. Minimizing the risk of using the new semiconductor process and keep the design schedule as planned were also the important design considerations. For this reason, simpler and robust circuit design is favored. Also a clocking scheme is chosen to reduce the work for timing closure by giving timing tuning flexibility to the designers.

Most of the design CAD tools are in-house developed. They are enhanced to handle the side effects of PD SOI technology, most notably, timing variation due to the floating body effect.

2. IMPLEMENTATION OVERVIEW

The 5th generation SPARC64 processor is fabricated in 130nm SOI CMOS with 8 layers of Cu metallization. It runs at 1.3GHz with 34.7W power dissipation in the laboratory. The chip contains 191M transistors with 19M in logic circuits, measures 18.14mm x 15.99mm and is covered with 5,858 low alpha emission lead bumps of which 269 are for I/O signals. The system bus is 16 bytes wide and operates with a 260MHz clock. It achieves a peak bandwidth of 4.16 GB/s in SDR and 8.32GB/s in DDR mode. Table 1 is the summary of the technology and the processor characteristics.

The processor core is a 4 issue out-of-order superscalar design with 11 stages of pipeline for basic integer instructions. Floating point add/multiply requires 5 more pipeline stages and load/store requires 4 more pipeline stages. The processor contains 2 way set associative level-1 instruction (L1I) and data (L1D) caches of 128KB each. The instruction fetch unit contains a 16K entry

Table 1. Implementation Summary

Technology	130nm CMOS SOI, Vdd=1.2V 8 Layer Cu metallization
Chip size	18.14mm x 15.99mm
Clock frequency	1.3GHz
Bus Interface	260MHz SDR/DDR, 16B wide
Power dissipation	34.7W (measured) @1.3GHz
Transistor count	191M (19M in Logic)
Micro architecture	4 issue out-of-order superscalar 11(~16) stage pipeline
Execution units	2 FX, 2 FP, 2 AGEN/LoadStore, 1 BR
L1 Cache	128KB 2way set associative I\$ and D\$, 16K entry 4way set associative BTAC
L2 Cache	2MB 4way set associative unified

(8B/entry) 4 way set associative branch target address cache (BTAC) with prediction information. A maximum of 4 instructions are issued into 4 groups of reservation stations. The execution unit contains 2 integer pipelines, 2 floating point pipelines, 2 load/store pipelines with address generator and one branch pipeline. The unified level-2 (L2) cache is 4 way set associative with 2MB capacity and achieves 41.6GB/s bandwidth at 1.3GHz clock. The chip microphotograph is shown in Figure 1. Microarchitecture descriptions of this processor are found in [2], [3], [4].




Figure 1. Chip Photomicrograph

3. HIGH RELIABILITY FEATURES

Highly reliable operation is an important requirement of this processor design as the CMOS scaling renders it more vulnerable to noise generated by alpha and cosmic particle hits, EMI and other

external noise, on top of internally generated noise. The L1D and L2 data, and L2 tag are protected by ECC. Tags of both L1I and L1D are parity protected and duplicated for correction. The L1I data, instruction TLB, data TLB and BTAC are parity protected. With these protections, the hardware can recover from a single bit error in these arrays without software assistance. In addition, this processor implements parity in all data path latches. About 80% of the 200K total latches are covered by parity check. The ALU and shifter have parity prediction circuits. Parity generated from the output signals is compared with the predicted one. The multiply/divide unit is protected with residue check [5] and parity prediction circuits. These parity bits are carried along the paths of the data flow and checked at the receiving ends. Figure 2 shows the error check mechanisms of this processor.

Instructions commit after error check. When an error is detected, the processor discards the intermediate states and re-issues only the instruction right after the last committed one. This scheme reduces internally generated noise for re-execution. Normal superscalar execution resumes after the completion of this instruction. The error will be corrected if the failure is intermittent in nature. This mechanism can also be a safety net for a signal integrity compromise due to the rare occurrence of a superposition of various noises.

4. TUNABLE SOFT BARRIER CLOCKING

The clock is distributed in a gridless tree structure with 4 levels of buffers from the PLL to the latch/flip-flops. Variable delay 3rd stage clock buffers and 4th stage clock choppers are used for clock tuning. Up to 500ps clock delay can be specified in 20ps increments at the time of the design. Combined with the relatively long 120ps transparent period of the pulsed latch, a large amount of time can be borrowed from the next pipeline stage. With this tunable soft barrier clocking scheme, critical path signals can travel multiple pipeline stages without being gated by the clock edge, minimizing the effect of clock skew and jitter. Figure 3 illustrates this scheme. The majority (85%) of the latches used clock without additional delay. The remaining 15% used up to 140ps delayed clock for balancing the delay between the adjacent pipe




Figure 2. Processor Block Diagram with Error Check Mechanisms

stages. All clock delays are back annotated based on parasitic extraction and the CAD tool reports set-up and hold violations.




Figure 3. Tunable Soft-barrier Clocking Scheme

5. CIRCUIT DESIGN

The standard cell library contains a total of 579 cells in both normal and low V_{th} transistor versions. Use of low V_{th} cells is restricted and results in 5.75% usage. The cell library contains 17 power levels for simple gates to facilitate delay tuning for long paths and race paths. Only 4 memory arrays and 67 macros ranging up to 12K transistors are designed using full custom circuits. The rest of the design is semi-custom, consisting of hand optimized assemblies of standard cells. All the circuits except RAMs are static. By limiting the amount of custom design and eliminating dynamic circuits, development time to the first tape out was reduced to 14 months. This short development time allowed the use of a more advanced semiconductor process which offsets the disadvantage in circuit speed and density compared to a highly custom approach which would require a longer development time.

6. POWER REDUCTION

Low power dissipation of 34.7W is attributable to: 1) static design, 2) clock and/or input gating of unused units to eliminate unnecessary toggles, 3) clock gating in the L2 cache unit to reduce clock power dissipation of 4Kx72bit RAM macros, 4) clock tree design which has lower capacitance than the big single-node clock grid, and 5) low parasitic capacitance 130nm CMOS SOI process.

7. DESIGN METHODOLOGY

Logic design is coded using Fujitsu's logic input language and then converted to Verilog. Logic designers generate unit level test vectors with the logic design of each unit. After the entire processor logic is integrated, those unit level test vectors are run on Fujitsu's LSP logic simulator to find logic bugs in each unit. Also the test vectors developed by the design verification team are run on the LSP logic simulator to verify the functionality of the whole processor.

When the processor logic becomes reasonably stable, CoBALT logic emulator [6] is used for massive logic simulation. An excess of 100B cycles of logic simulation have been performed to ensure functional correctness of the design.

The processor chip is divided into sub-chips which corresponds to high level logic functions, like instruction unit, execution unit etc.

and chip level structure, like external I/O, fuse etc.. These sub-chips are mostly rectangular in physical shape and connections between them are made by abutment.

These sub-chips are populated by layout sub-groups (LSGs). The placement of LSGs within a sub-chip is manual, also the critical nets are routed manually. Routing channels used for these critical path nets are communicated as a keep out information to the lower level intra-LSG routing.

Those LSGs are made of pseudo macros and library cells. A pseudo macro is a hard macro made from the pre-placed library cells and manually pre-routed nets. Placement of small cells and wiring within an LSG is mostly manual with the assist of in-house place and route tool. This physical design hierarchy is shown in Figure 4.




Figure 4. Physical Design Hierarchy




Figure 5. Hold Violation Fix Advise

A timing of the chip is calculated by the in-house static timing analyzer based on the logic data base, physical design data base, library cell timing files and other technology files. The analyzer not only reports setup and hold violations at the receiving latch inputs, it has the capability to report slacks for each input of the gates within the logic cone leading to the violating input as shown

in Figure 5. This information helps logic designers to identify where to improve timing and/or the right place to insert delay buffer to eliminate hold violation. This functionality of the timing analyzer is very useful since the pulsed latch based design with wide transparency window tends to increase the number of hold time violations.

The static timing analyzer is enhanced to take the delay variation of the PD SOI circuit into calculation. Also the new tool is developed to generate timing models of custom circuits with the floating body effect considered.

8. CONCLUSIONS

The 5th generation SPARC64 processor achieved 1.3GHz clock with 34.7W power dissipation (current Fujitsu unix server products use slightly higher Vdd to run this processor at 1.35GHz and resulted power dissipation is about 45W). This power consumption is much smaller compare to the one of the other high end microprocessors. The processor implements error detection mechanism for the execution units and data paths and recovers from the detected errors by instruction retry. We estimate that this retry improves the MTBF of the processor by an order of magnitude. The implementation of this level of high reliability features is industry first for the RISC processors.

Simpler circuit design and tunable clocking scheme contribute to achieve short development time of 14 months from the start of the implementation to the first tape out.

9. ACKNOWLEDGEMENTS

The authors acknowledge the contributions from the rest of the design team. Special thanks to Akira Kaneko, Noriyuki Itoh, Takeshi Ibusuki, Kimihiro Suzuki for their leadership.

10. REFERENCES

- [1] H.Ando, Y.Yoshida *et al* "A 1.3GHz Fifth Generation SPARC64 Microprocessor", ISSCC Dig. Tech Papers, Feb. 2003, pp.246-247.
- [2] A.Inoue "Fujitsu's New SPARC64 V for Mission-Critical Servers", Microprocessor Forum 2002, Oct. 2002.
- [3] A.Inoue "SPARC64 V for Unix Servers" (in Japanese), Fujitsu Vol 53.6, Nov. 2002, pp.450-455.
- [4] M.Sakamoto, A Katsuno *et al* "Microarchitecture and Performance Analysis of a SPARC-V9 Microprocessor for Enterprise Server systems", Proceedings of HPCA9, Feb. 2003, pp.141-152.
- [5] J.Watterson, J.Hallenbeck, "Modulo 3 residue checker: a new results on performance and cost", IEEE Trans. Computers, Vol. 37, May 1988, pp.608-612.
- [6] http://www.quickturn.com/products/cobaltplus_data_sheet.htm