2004 IEEE International Conference on Integrated Circuit Design and Technology

Deep-submicron Design Challenges for a Dual-Core 64b UltraSPARC Microprocessor Implementation

Toshinari Takayanagi, Jinuk Luke Shin, Jeffrey Su and Ana Sonia Leon
Sun Microsystems, Inc.
430 N. Mary Ave. Sunnyvale, CA 94085, U.S.A.
+1-408-720-4894
{toshinari.takayanagi, jinuk.shin, jeffrey.su, ana-sonia.leon}@sun.com

Abstract

A processor core, originally designed in a 0.5µm Al process, is redesigned for a 0.13µm Cu process to create a dual-core processor with 1MB integrated L2 cache, offering an efficient performance to power ratio for compute-dense sever applications. Circuit design challenges, including negative bias temperature instability (NBTI), leakage, coupling noise and intra die process variation are discussed.

1. Introduction

Three factors are driving a paradigm shift in 64b microprocessor design for enterprise servers, from single-thread processors based on a complex wide-issue/deep-pipeline core to multi-thread designs based on multiple simpler cores. The first factor is that efforts to elaborate core designs in order to increase clock frequency and instructions-per-cycle (IPC) are both reaching the point of diminishing returns [1]. The added core complexity has also increased the design cycle enormously. The second factor is sharply increasing power consumption, which has made the ability to cool systems a serious concern. The third factor is the change in the nature of application workloads in network computing. For network computing applications, total throughput or the number of threads executed per unit of time is more important than decreasing the amount of time needed to execute a single-thread. Network-based applications also have rich thread-level parallelism which multi-threaded hardware can effectively exploit.

In light of the above considerations, a project to create a multi-thread processor specialized for low-cost network servers for compute-dense environments, such as rack-mount and blade system, is based on reusing the relatively simple and power-efficient UltraSPARC II core in a dual-core configuration [2][3]. Not only did the UltraSPARC II core provide an optimum solution for this specific design point but reusing this proven core also helped to achieve a short design cycle.

Porting the UltraSPARC II core, which was originally implemented in 0.5µm/3.3V technology, and last implemented in 0.25µm/2.1V technology, to 0.13µm/1.3V technology raised deep-submicron design challenges facing the semiconductor industry today; Negative bias temperature instability (NBTI), coupling noise, leakage current and intra-die process variation have all impaired circuit margins, scalability and reliability. Meeting these challenges required stringent circuit analysis and re-design, as discussed in this paper.

2. Chip Overview

This 64b SPARC processor is designed for compute-dense systems such as rack-mount and blade servers for network computing, as is described above. The critical requirements for this type of applications are high computing throughput, high performance-per-watt ratio, enterprise-class 64b memory subsystem and reliability. A short design cycle was also critical for this project.

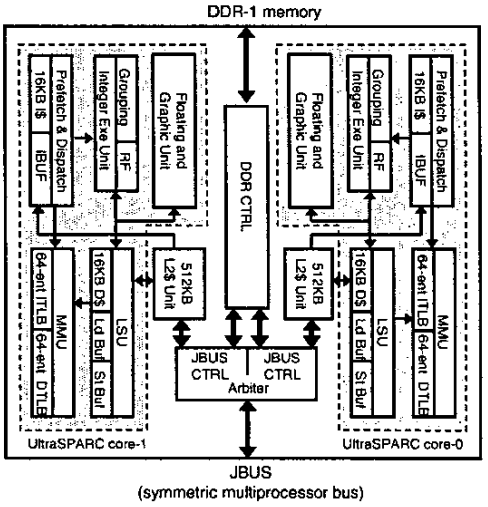

Figure 1. Chip block diagram

The chip contains two UltraSPARC II cores, two 512KB L2 cache subsystems, DDR-1 memory controller and symmetric multiprocessor bus (JBus) interfaces (Figure 1). The core provides an efficient performance-per-watt having balanced hardware complexity with 4-issue superscalar, 9-stage pipeline and in-order execution/out-of-order completion [4]. The core includes 16KB 2-way set-

associative I-cache, 16KB directly-mapped D-cache, 2KB next-field RAM, 64-entry full associative I-TLB and D-TLB. Each L2 cache is 4-way set-associative.

Parity bits are added to L1 cache data/tag and L2 tag arrays while L2 data arrays are protected by error correction code (ECC). The memory controller supports up to 16GB of physical memory. JBus controllers allow low-cost multiprocessing systems with configurations of up to four-chips without any glue logic. Typical power dissipation at 1.2GHz and 1.3V is 23Watts, which is the lowest published figure for 64b server processors [2].

The chip is fabricated in Texas Instruments' advanced 0.13μm CMOS process with 7 layers of Cu and a low-k dielectric. The chip micrograph is shown in Figure 2. The transistor count is 80M, of which 72M is SRAM. The 206mm^2 die is packaged in a 959-pin ceramic μPGA. The chip interface is made pin-compatible with UltraSPARC IIIi processor [5] to effectively reuse the existing system resources.

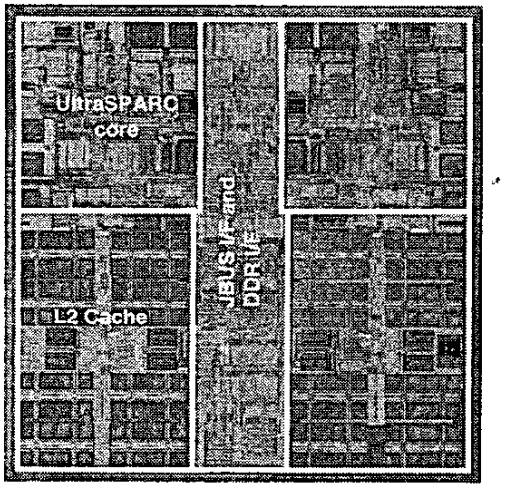

Figure 2. Chip micrograph

3. Deep-Submicron Design Challenges

3.1 NBTI

NBTI is the aging effect that decreases PMOS current mainly due to Vt shift over silicon lifetime [7]. This Vt shift is strongly dependent on gate-source bias and temperature but barely dependent on drain voltage. The NBTI impact on the circuits includes speed degradation, increased delay variation, shifted PMOS/NMOS drive current ratio, decreased Vdd/Vt headroom and Vt mismatch. Many circuits were modified to enhance margins for NBTI.

Particularly, current-mode latch sense amps used for L1 caches and TLBs (Figure 3-a)[6] were highly affected, degrading the total sense delay by 42% (Figure 4-c). The cross-coupled PMOS pair, M3 and M4, which act as low input impedance devices during equilibration and positive feedback while sensing, were unevenly affected due to unequal 0 or 1 read rate. This situation is particularly worse for TLB matchline sense amp as each entry of the TLB will read out the "miss" data most of the time while one of the other 63 entries provides the "hit" signal. This could cause a worst case Vt mismatch of about 50mV between the PMOS pair, requiring a longer signal development time to overcome the offset. The PMOS Vt shift also attenuated the gain.

To cope with NBTI, the cross-coupled PMOS pair M3 and M4 are replaced by T3 and T4 whose gates are commonly biased at about 40% of Vdd during sensing. These act as low impedance devices being biased in saturation mode. As the gate bias is identical for T3 and T4, the Vt imbalance is minimized. In addition, PMOS T1 and T2 are added to speed up the low-to-high output transition. Although T1 and T2 can get an uneven Vt shift, it is not critical as they get activated after the significant part of the amplification is completed. These modifications improved the deteriorated sense delay by 22%, achieving 15% speedup over voltage sense amps (Figure 4-d).

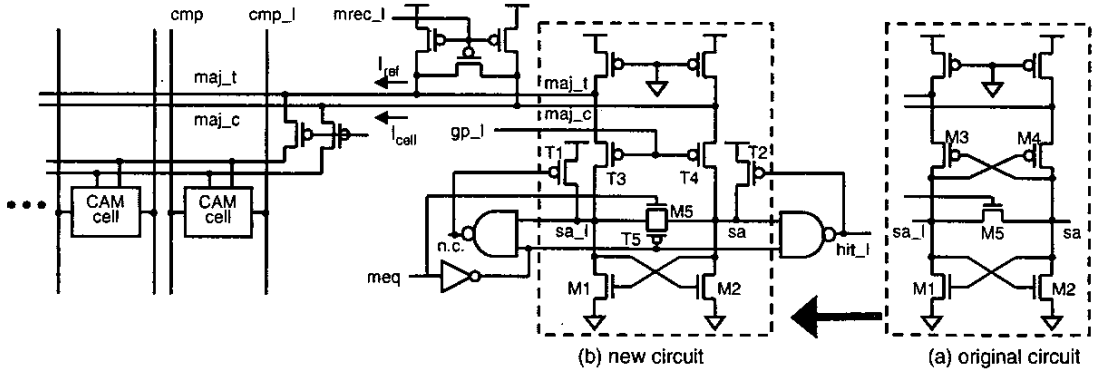

Figure 3. TLB matchline sense amplifier

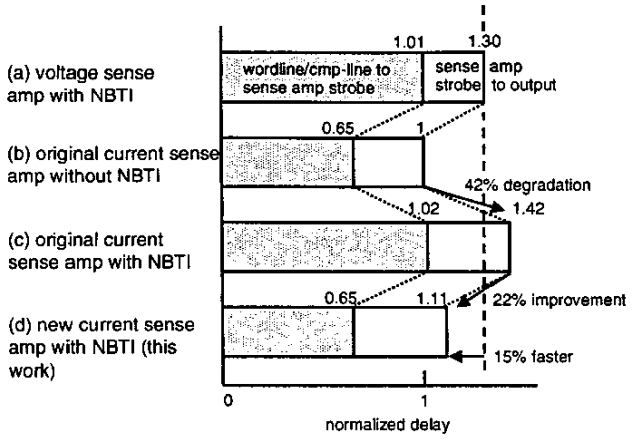

Figure 4. Sense amplifier delay comparison

3.2 Coupling Noise and Leakage Current

An in-house CAD tool described in [8] with numerous enhancements was applied to this design to exhaustively analyze the coupling noise in interconnects. As the tool first performs a worst-case analysis, assuming all the aggressors switch simultaneously, the reported results were inherently conservative. Accounting for the logic and timing constraints of the victim and aggressor nets effectively filtered out the false violations.

Increased leakage current made dynamic gates and open-drain pass-gate latches/muxes very noise-sensitive structures. In order to address leakage and noise issues, additional circuit modifications were required. The I-cache wordline detector in Figure 5 is one example.

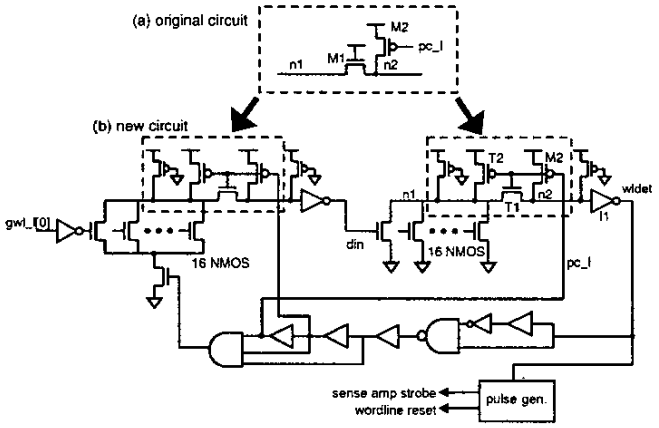

Figure 5. I-cache wordline detector

This circuit is a 256-inputs OR gate, which consists of two levels of 16-inputs self-resetting dynamic OR gates, to detect a wordline transition for sense amp strobe. The wired-NOR net, n1, is susceptible to leakage and noise, as 16 NMOSs are connected in parallel with a long wire. In the original circuit (Figure 5-a), n1 was precharged to Vdd-Vt for speed, however the noise margin of this circuit was reduced due to the lower supply voltage: a 100mV drop at n1 could cause significant output glitch, as M1 turns on easily to discharge node n2. The simulation waveforms are shown in Figure 6.

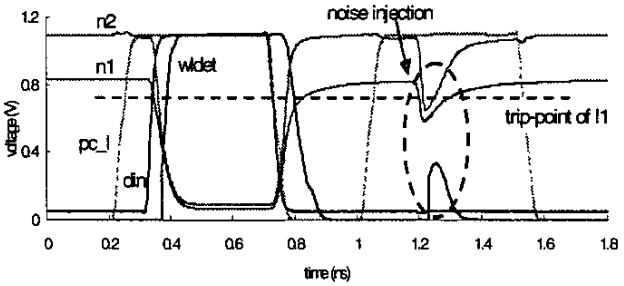

Figure 6. Simulated waveforms of original wordline-detector with noise injection

In the new circuit (Figure 5-b), n1 and n2 are both precharged to Vdd with NMOS T1 between them. The gate of T1 is high during the evaluation and low during the precharge. During the evaluation, T1 acts as a noise decoupler since the voltage drop at n1 does not propagate to n2 unless it is large enough to turn on T1. In addition, T1 decouples n2 from the large capacitance on n1 during the precharge, speeding up the reset path. Compared to a conventional domino gate, this circuit achieves similar speed while improving the noise margin by 35%. The wordline detection slowed down by 9% from the original circuit, but the cache access time is not impacted since the extra delay is absorbed in the sense amp strobe buffering stage.

3.3 Intra-die Process Variation

Since clock skew does not scale proportionally to gate delay due to intra-die process variation, a significant number of new hold-time violations had to be addressed. These needed to be fixed with minimal physical design changes for fast area- and timing-closure. New hold-time hardened flops that have either larger output delay or smaller hold-time while keeping the same footprints were created for the most frequently used flops. An automated flow was developed to identify and replace the flops to fix the hold-time violation without affecting the critical paths. More than 60% of the total hold-time violations were solved by this methodology.

One challenge here was to create a larger output delay while keeping the original flop size. The flop depicted in Figure 7 utilizes the scan slave path for normal output. In this circuit topology and clocking scheme, the pass-gate

TG1 is controlled normally-on in the operation mode with se=0 and thus sclk=1. This achieves an additional three-stage gate delay without increasing the flop footprint.

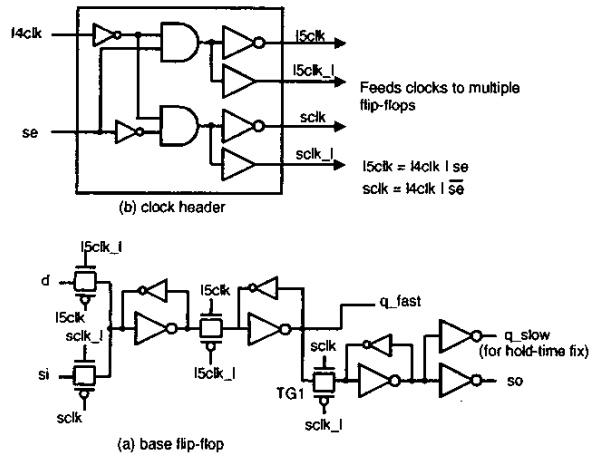

Figure 7. Scan flip-flop with delayed output

Intra-die process variation also exhibits increasing challenges for designing SRAM sense amps and memory cells. While the main design methodology to check the circuit tolerances was to add worst-case intra-die process offsets to the corner simulations, this methodology was supplemented with a number of statistical simulations to make the analysis more realistic.

4. Conclusion

Deep-submicron design challenges encountered in porting a processor core originally designed in 0.5µm technology to contemporary 0.13µm technology have been discussed. Circuits that required redesign to cope with NBTI, coupling noise, leakage current and intra-die process variation were presented in detail.

Together with the integrated DDR-1 memory controller, JBus interfaces and newly designed high-performance 1MB L2 cache subsystems, a dual-core 64b microprocessor optimized for low-cost network servers for compute-dense environments is successfully created. The target frequency of 1.2GHz at 1.3V, 85°C is achieved with comfortable margin (Figure 8), dissipating only 23W with typical application workloads. Reuse of the UltraSPARC II core contributed to the chip's low power consumption, and also helped achieve a short design cycle for this project.

5. Acknowledgement

The authors acknowledge contributions from the entire Gemini development team and all those who supported the project both at Sun Microsystems and Texas Instruments. Special thanks to Raymond Heald, James Kaku and Harlan McGhan for their technical advice, and Andre Kowalczyk, Renu Raman, Allan Strong, Fred DeSantis, Hunter Ward and Julie England for supporting the project.

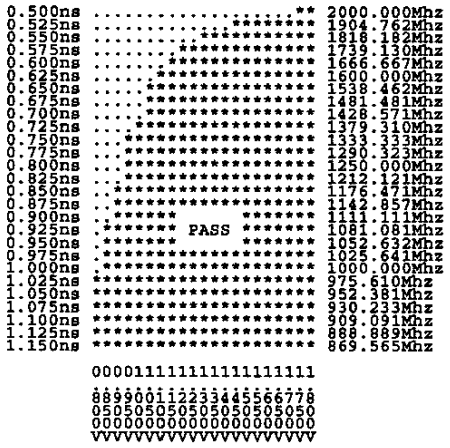

Figure 8. Shmoo plot at 85°C

6. References

[1] M. Horowitz and W. Dally, "How Scaling Will Change Processor Architecture," ISSCC Dig. Tech. Papers, Feb. 2004, pp. 132-133

[2] S. Kapil, "Gemini: A Power-efficient Chip Multi-Threaded UltraSPARC Processor," 15th Hot Chips Symposium, Aug. 2003

[3] T. Takayanagi et al., "A Dual-Core 64b UltraSPARC Microprocessor for Dense Server Applications," ISSCC Dig. Tech. Papers, Feb. 2004, pp. 58-59

[4] Lev, L. A., et al, "A 64b microprocessor with multimedia support," IEEE J. Solid State Circuits, vol. 30, no. 11, Nov. 1995, pp. 1227-1238

[5] G. K. Konstadinidis et al., "Implementation of a Third-Generation 1.1-GHz 64-bit Microprocessor," IEEE J. Solid State Circuits, vol. 37, no. 11, Nov. 2002

[6] E. Anderson, "A 64-Entry 167MHz Fully-Associative TLB for a RISC Microprocessor," ISSCC Dig. Tech. Papers, Feb. 1996, pp. 360-361

[7] V. Reddy et al., "Impact of Negative Bias Temperature Instability of Digital Circuit Reliability," Reliability Physics Symposium, 2002, pp. 248-254

[8] K. Aingaran et al., "Coupling Noise Analysis for VLSI and ULSI Circuits," ISQED 2000 proceedings, March 2000, pp. 485-489