

A Dual-Core 64-bit UltraSPARC Microprocessor for Dense Server Applications

Toshinari Takayanagi, *Member, IEEE*, Jinuk Luke Shin, *Member, IEEE*, Bruce Petrick, Jeffrey Y. Su, Howard Levy, Ha Pham, Jinseung Son, Nathan Moon, Dina Bistry, Umesh Nair, Mandeep Singh, Vikas Mathur, and Ana Sonia Leon, *Member, IEEE*

Abstract—A dual-core 64-bit microprocessor optimized for compute-dense systems such as rack-mount and blade servers for network computing was developed. The chip consists of two UltraSPARC II cores, each with its own 512 kB L2 cache, a DDR-1 memory controller, and symmetric multiprocessor bus (JBus) controllers. The 206-mm² die is fabricated in 0.13- μ m CMOS technology with seven layers of Cu and a low- k dielectric. The chip offers a highly efficient performance-per-watt ratio with a typical power dissipation of 23 W at 1.3 V and 1.2 GHz. A short design cycle was achieved by leveraging existing designs wherever possible and developing effective design methodologies and flows. Significant design challenges faced by this project are described. These include deep-submicron design issues, such as negative bias temperature instability (NBTI), leakage, coupling noise, intra-die process variation, and electromigration (EM). A second important design challenge was implementing a high-performance L2 cache subsystem with a short four-cycle core-to-L2 latency including ECC.

Index Terms—Chip Multithreading (CMT), coupling noise, current-mode sense amplifier, deep-submicron technology, dense server, dual-core, ECC, electromigration, hold time, leakage, L2 Cache, microprocessor, multicore, multiprocessor, multithread, negative bias temperature instability (NBTI), process variation, thread-level parallelism (TLP), translation look aside buffer (TLB), UltraSPARC.

I. INTRODUCTION

FOR THE PAST two decades, microprocessor performance has doubled every 18–24 months. This level of performance scaling has been sustained by increasing both clock frequency and instruction-level parallelism (ILP). However, these two factors are now reaching the point of diminishing returns [1]. Typical application programs have only limited ILP. Deep-pipelining, wide superscalar, out-of-order, and speculative execution have increased silicon area, power dissipation, design complexity, development cost, and project time enormously, compared to the actual performance gains realized by implementing these features. The ever-increasing gap between processor and memory frequency also hinders overall system throughput.

In addition, network computing based on today's pervasive use of the Internet has drastically changed the nature of application workloads. Network-based applications, such as online transaction processing, are rich in thread-level parallelism (TLP). TLP-rich applications require high computing

throughput to execute multiple threads/processes simultaneously rather than high single thread performance.

Furthermore, power consumption has become a critical issue. Power consumption of complex processors today is often over 100 W, making the cost of system cooling a serious concern. This is particularly true for dense server environments in the data center in which compute density is vital to efficiently utilize expensive space. In power-constrained environments, energy efficiency is more important than the highest possible performance.

This paper describes a 64-bit UltraSPARC processor explicitly designed to address the above points. The major goals were to create a low-cost processor that provides high computing throughput for today's network-centric applications with low power dissipation. A secondary goal was to minimize development time and cost. The solution adopted was to create a highly integrated multicore design, based on an existing relatively simple processor core, with on-chip L2 cache subsystems and peripheral controllers.

The discussion focuses on two of the major challenges encountered in this development effort. The first challenge was to port the selected UltraSPARC II core design, last implemented in a 0.25 μ m/1.9 V process, to a current 0.13 μ m/1.3 V process, coping with the deep-submicron technology difficulties facing the semiconductor industry today. The second challenge was to design a new on-chip high-performance, high-capacity, high-reliability L2 cache subsystem, replacing the off-chip L2 cache used by earlier UltraSPARC II processors.

Section II provides an overview of the design targets, solutions and the chip microarchitecture. Section III describes solutions to the issues raised by migrating the older implementation to current deep submicron technology. Section IV details the L2 cache subsystem design. Section V briefly describes the chip integration with some statistical data. Finally, Section VI provides measured results from silicon and conclusions.

II. ARCHITECTURE AND CHIP OVERVIEW

This 64-bit SPARC processor is designed for compute-dense systems such as rack-mount and blade servers for network computing [2], [3]. Critical requirements for these types of applications are high compute throughput, high memory bandwidth, a large addressing space, high reliability, low power, and low cost. A short design cycle was also critical for this project.

To address the above design targets, the optimum solution is to create an on-chip dual-core processor based on the UltraSPARC I/II microarchitecture [4], [5] with embedded 1 MB

Manuscript received April 15, 2004; revised June 30, 2004.

The authors are with Sun Microsystems Inc., Sunnyvale, CA 94085 USA (e-mail: jinuk.shin@sun.com; ana-sonia.leon@sun.com).

Digital Object Identifier 10.1109/JSSC.2004.838023

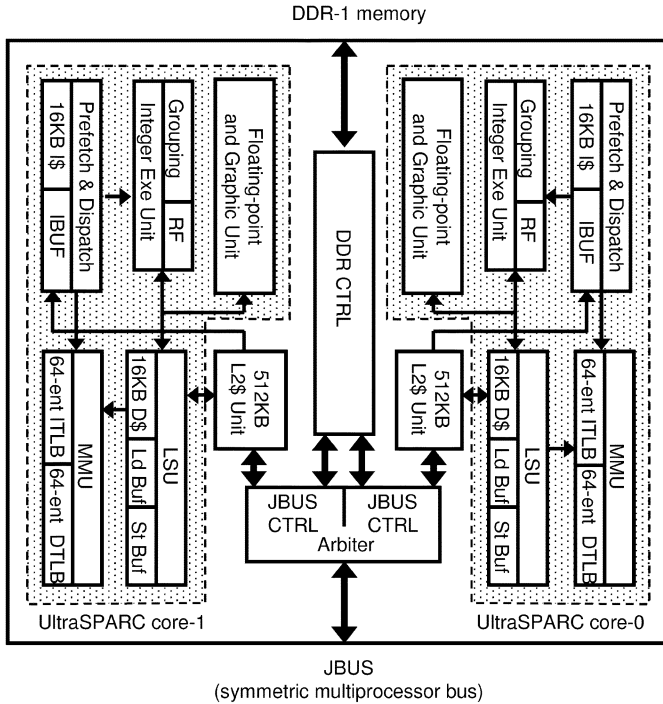


Fig. 1. Block diagram.




Fig. 2. Chip microphotograph.

L2 cache, DDR-1 memory controller and symmetric multiprocessor bus (JBUS) interfaces (Figs. 1 and 2). This core achieves efficient performance-per-watt with a minimum degree of hardware complexity. Major performance features include a 4-issue superscalar instruction dispatch, a 9-stage pipeline (Fig. 3) and in-order execution/out-of-order completion.

To help minimize power consumption, the implementation of this core relies on a predominantly static circuit design style and a balanced H-tree clock distribution. Typical power dissipation at 1.2 GHz and 1.3 V is 23 W. Breakdown of the

power dissipation for the chip and core is shown in Fig. 4. Measured performance-to-power ratio for the SPECint_rate2000 throughput benchmark is 0.45 SPECint_rate2000/W (preliminary data). This is one of the most efficient ratios published to date for 64-bit server processors [2].

Parity bits are added to the L1 cache data/tag and L2 tag arrays while L2 data arrays are protected by ECC for enterprise-class data reliability. The memory controller supports up to 16 GB of physical memory. The JBus controllers allow low-cost multiprocessing systems with configurations of up to four chips (eight threads) without any glue logic. The chip implements system software interfaces to manage the multiple threads on a die in accordance with standards developed by Sun to support its broad family of forthcoming chip multithreading (CMT) designs.

The chip is fabricated in Texas Instruments' advanced 0.13- μm CMOS process with seven layers of Cu and a low- k dielectric. The transistor count is 80 M, of which 72 M is SRAM. The 206 mm² die is packaged in a 959-pin ceramic μPGA . The chip interface is made pin-compatible with the UltraSPARC IIIi processor [6] to effectively reuse existing system resources. The core and chip features are summarized in Table I with the corresponding features of the UltraSPARC II processor.

The core circuits were originally implemented in 0.5 μm /3.3 V and last implemented in 0.25 μm /1.9 V technology. Redesigning these circuits for 0.13 μm /1.3 V technology raised typical deep-submicron design challenges, including negative bias temperature instability (NBTI), leakage, coupling noise, intra-die process variation, and electromigration (EM). The project goals also required achieving a short on-chip L2 cache latency with ECC. These challenges are discussed in detail in the next two sections.

III. DEEP-SUBMICRON DESIGN CHALLENGES

A. Circuit Design Methodologies

The original UltraSPARC I was designed with great emphasis on future technology migrations, applying twelve simulation corners coupled with a concept of "speed-up ratio" to check circuit scalability [4]. Many circuits scaled well from the original 0.5- μm technology to the current 0.13- μm technology thanks to this robust methodology. However, new deep-submicron technology issues such as increased transistor leakage, not anticipated at the time of the original design, impaired circuit scalability. Additional circuit design guidelines and new rule checkers were developed and put in place, including a rebalanced keeper ratio for dynamic gates, latch sizing rules for both writability and stability, and length limits for wires connecting to pass gate diffusions. The twelve simulation corners, which were defined in terms of PMOS, NMOS, supply voltage and temperature conditions, were extended to 24 corners, including high-leakage and ultra-high/low voltage conditions for circuit stability and reliability. As an example, the original circuit, with alternating P and N dynamic stages (Fig. 5), failed both the circuit guideline checker and high-leakage corner simulation, requiring upsizing of the keeper devices. The fast corner simulations also showed that the precharge pulses of the original

Integer

F Fetch	D Decode	G Group	E Execute	C Cache Access	N1	N2	N3	W Write
Instructions are fetched from I-cache	Instructions are decoded and placed in the instruction buffer	Up to 4 instructions are grouped and dispatched. IRF accessed	Integer instructions are executed and virtual address calculated	Dcache/TLB accessed. Branch resolved	Dcache hit or miss determined. Deferred load enters load buffer	Integer pipe waits for floating-point/graphic pipe	Traps are resolved	All results are written to the register files. Instructions are committed

R Register	X1	X2	X3
Floating-point/graphics instructions are further decoded. FRF accessed	Start execution	Execution continued	Finish execution

Floating-point/Graphics

Fig. 3. Instruction pipelines.




Fig. 4. Chip and core power dissipation breakdown. (a) Chip power: 23 W (typical). (b) Processor core power: 5.3 W (typical).





Fig. 5. Dynamic gate circuits were analyzed by using circuit checkers and high-leakage corner simulation.

self-resetting gates were insufficient, requiring the insertion of two additional inverters (Fig. 6).

The circuit design methodology includes self-timed margins (STMs) between cycle-time-independent racing paths, depicted in Fig. 7, as a quantitative measure of circuit tolerance for variations in process, voltage, and temperature. For example, an STM of 20% means that the timing relation of two signals is satisfied

Fig. 6. Delay addition to self-resetting gate. Due to gate speedup in the 0.13- μ m process, the original self-resetting pulse did not reach full ground at fast corner simulation.

even if path D1 slows down by 20% and path D2 speeds up by 20%.

TABLE I
CHIP FEATURE SUMMARY WITH COMPARISON TO ULTRASPARC II

	UltraSPARC II (1999)	Dual-Core UltraSPARC (2003)
Architecture	64b SPARC-V9 with multimedia instruction extensions	<ul style="list-style-type: none"> • Dual-core with CMT features (core disable/parking/error reporting)
Pipeline	<ul style="list-style-type: none"> • 4-issue superscalar (2 integer ops, 2 FPU/graphics ops/cycle) • 9-stage pipeline • In-order execution/out-of-order completion 	<ul style="list-style-type: none"> • Core is based on UltraSPARC II microarchitecture • DDR-1 memory controller • JBUS (SMP bus) interfaces • E*Star mode support • L1 caches are parity protected
L1 cache/MMU	<ul style="list-style-type: none"> • 16KB 2-way I-cache • 2KB next field RAM • 16KB direct D-cache • 64-entry full-associative I-TLB and D-TLB 	
L2 cache	Support off-chip direct-map 256KB – 16MB L2 cache	<ul style="list-style-type: none"> • On-chip unified 512KB 4-way set-associative cache for each core • ECC protected • 4-cycle latency to core / 9-cycle load-use latency • 2-cycle data throughput
Process	0.25 μ m CMOS/5 layers Al (Original UltraSPARC I was fabricated in 0.5 μ m CMOS/4 layers Al.)	0.13 μ m CMOS/7 layers Cu
Voltage	1.9V	1.3V
Transistors	5.4M	80M total, 72M for SRAMs
Die size	126mm ²	206mm ² , 29 mm ² per core
Clock freq.	450MHz	1.2GHz
Power	23W	23W (typical), 5.3W per core
Max. memory bandwidth	3.6GB/s to L2 1.92GB/s to main memory	9.6GB/s to L2 4.26GB/s to main memory
Performance	2.5 SPECint_rate2000 (estimated)	11.0 SPECint_rate2000 (preliminary)

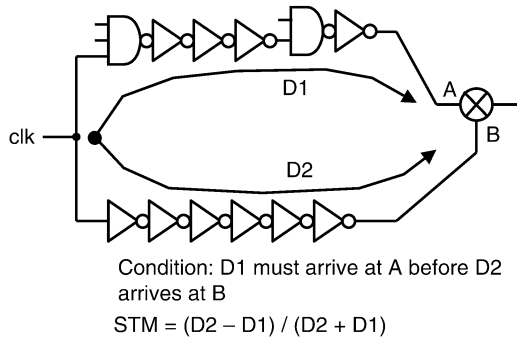


Fig. 7. Self-timed margin (STM) is a measure defined to evaluate delay variation tolerance of racing paths. For example, memory bitline signal is D1 and sense amp strobe is D2 for the sense amp evaluation.

B. NBTI

NBTI is the aging effect that decreases PMOS current mainly due to V_t shift over the silicon lifetime [10], [11]. This V_t shift is strongly dependent on gate-source bias and temperature but barely dependent on drain voltage. The damage mechanism is related to hole trapping in the gate oxide and interface state generation. The impact of NBTI includes speed degradation, increased delay variation, shifted PMOS/NMOS drive current ratio, decreased V_{dd}/V_t headroom and increased V_t mismatch. To accommodate additional delay variations anticipated due to NBTI, the STM requirement was raised appropriately. Since the

delay variation caused by NBTI is unidirectional, only positive variation factors needed to be considered. For example, if the worst speed degradation due to NBTI is 4%, the STM requirement is raised by 2%.

In particular, current-mode latch sense amps used for L1 caches and translation lookaside buffers (TLBs) [Fig. 8(a)] [7] were highly affected by NBTI, degrading the total sense delay by 42% [see Fig. 10(c)]. The cross-coupled PMOS pair M3 and M4 which act as low input impedance devices during equilibration [8], [9] and positive feedback while sensing, were unevenly affected due to unequal 0 or 1 read rate. This situation is particularly bad for the 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 V_t mismatch of about 50 mV between the PMOS pair, requiring a longer signal development time to overcome the offset. The PMOS V_t 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 controlled by gp.l [Fig. 8(b)]. The sense amp operation is as follows (Fig. 9). The CAM search operation is initiated from the clock rising edge. The CAM cells start sinking current with mismatch, and meq is activated to equilibrate the sense amp. PMOS T5 is added to effectively equalize sa and sa.l nodes at the lowered supply voltage. Then gp.l is lowered to about 40% V_{dd} to

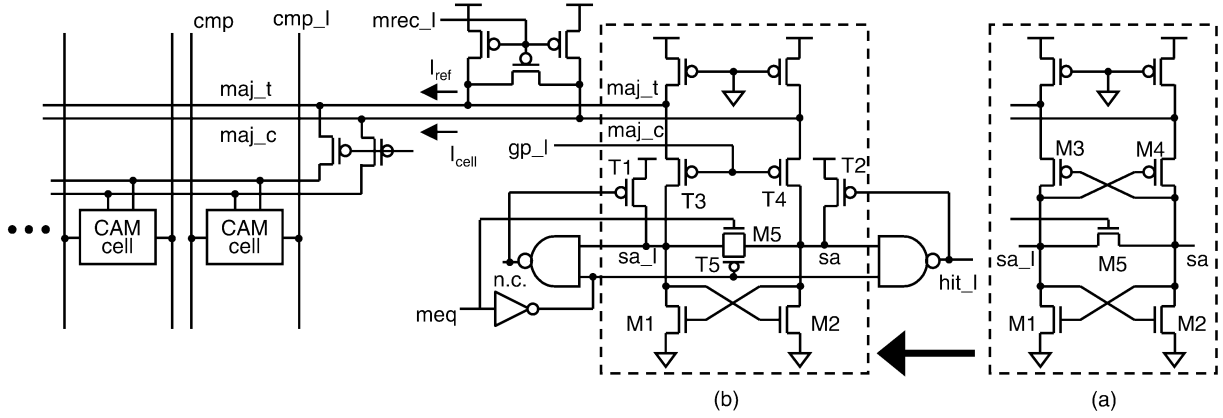


Fig. 8. TLB matchline sense amplifier. (a) Original circuit. (b) New circuit.

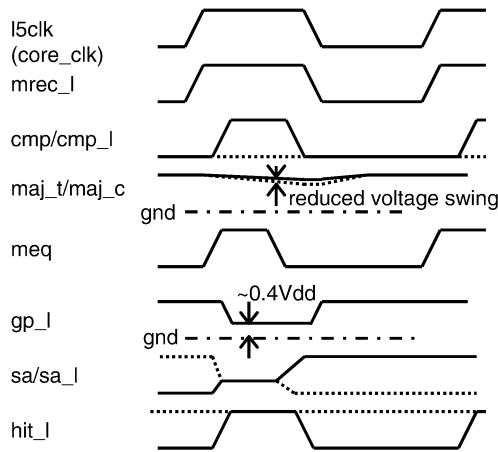


Fig. 9. Timing diagram of TLB sense amp operation.

bias the PMOS pair T3 and T4 in saturation mode to act as low impedance devices for current-mode sensing. After a sufficient signal develops to overcome the offsets of the matchline and sense amp, meq is taken low to kick off positive feedback amplification. This positive amplification is essentially left to the NMOS pair M1 and M2. The NAND gates are activated after one gate delay from meq to detect the transition of sa and sa_l. Once either sa or sa_l crosses the trip point of the nand gate, the newly added T1 or T2 turns on to accelerate the low-to-high output transition. After the sense action is done, gp_l is raised to V_{dd} to decouple the matchline from the sense amp, allowing the matchline to be precharged back to V_{dd} .

As the gate bias is identical for T3 and T4, the V_t imbalance is minimized. Although T1 and T2 can get an uneven V_t 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 a 15% speedup over the voltage sense amps [Fig. 10(d)].

C. Leakage and Coupling Noise

In order to address leakage and noise issues, additional circuit modifications are required. The I-cache wordline detector in Fig. 11 is one example. This circuit is a 256-input OR gate, which consists of two levels of 16-input self-resetting dynamic




Fig. 10. Delay comparison between conventional voltage-mode sense amp and the current-mode sense amps with and without NBTI effect. (a) Voltage sense amp with NBTI. (b) Original current sense amp without NBTI. (c) original current sense amp with NBTI. (d) New current sense amp with NBTI (this work).

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, n1 was precharged to $V_{dd} - V_t$ for speed, however the noise margin of this circuit was reduced due to the lower supply voltage: a 100-mV drop at n1 could cause the circuit to fail as M1 turns on easily to discharge node n2.

In the new circuit, n1 and n2 are both precharged to V_{dd} 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%. For this case, noise margin is defined as an allowable noise voltage at n1 that causes an output glitch of 10% of V_{dd} at n3. Circuit optimization nullified the speed impact caused by the on-resistance of T1 and the

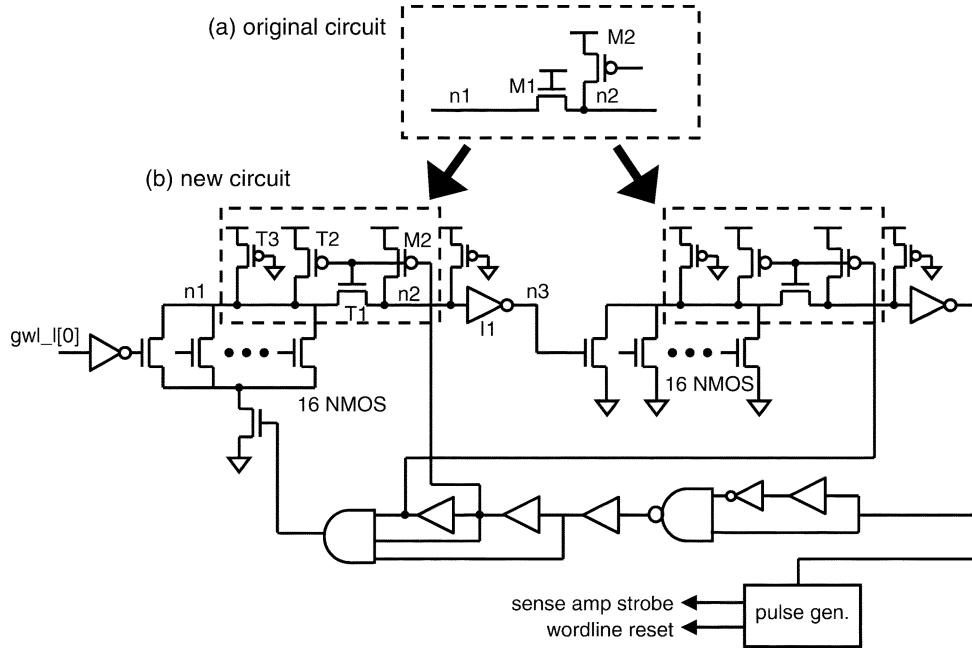


Fig. 11. I-cache wordline detector.

newly added keeper T3. While the wordline detection slowed down by 9% from the original circuit, the cache access time is not impacted since the extra delay is absorbed in the sense amp strobe buffering stage.

D. 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 are created (Fig. 12). Fixing hold-time violations is a challenging problem with respect to both timing closure and silicon area, as inserting buffers into the hold-time violation paths can affect critical paths and increase area. In addition, it is an algorithmically complex problem, since many possible solutions exist in terms of where delays could be inserted and how many violations can be solved at a time. To effectively solve hold-time violations, a new automated flow was created, relying on various delay buffers and footprint compatible hold-time hardened flops with either larger output delay or smaller hold-time.

Replacing flops with corresponding footprint compatible hold-time hardened flops is an efficient way of solving timing violations, as it does not affect the placement and routing of the block. 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 Fig. 13 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.

A chart of the hold-time fix flow with a hypothetical example is shown in Fig. 14. First, the flow identifies the timing requirements for hold-time and critical path at each node for the violated paths network based on the timing analysis results. Second, it performs a depth-first-search (DFS) on the network




Fig. 12. Hold-time violation distribution of the processor core. The Y-axis shows hold-time violation in ps between two flops, while the X-axis shows minimum critical path slack in which either the source or destination flop is involved. The simple flop swap method is applicable when the critical path slack is larger than the hold-time violation.

to identify substantially all possible solution scenarios with various numbers of delay buffers to be inserted at various nodes. Third, the optimum solution out of this set is determined by taking two factors into account: total number of possible fixes and the fixes-per-buffer ratio. The algorithm selects a solution scenario with the largest fixes-per-buffer ratio amongst those having a number of fixes greater than the mean of the set. Next, the flow inserts actual gates into the netlist. Delay calculations up to this point are done based on a linear RC model for runtime speed reasons. Then, accurate critical path timing analysis is performed again and any buffers causing critical path violations are removed at this stage. Iterative runs based on this algorithm left only a minimum number of fixes to be done manually.

Intra-die process variation also causes 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.

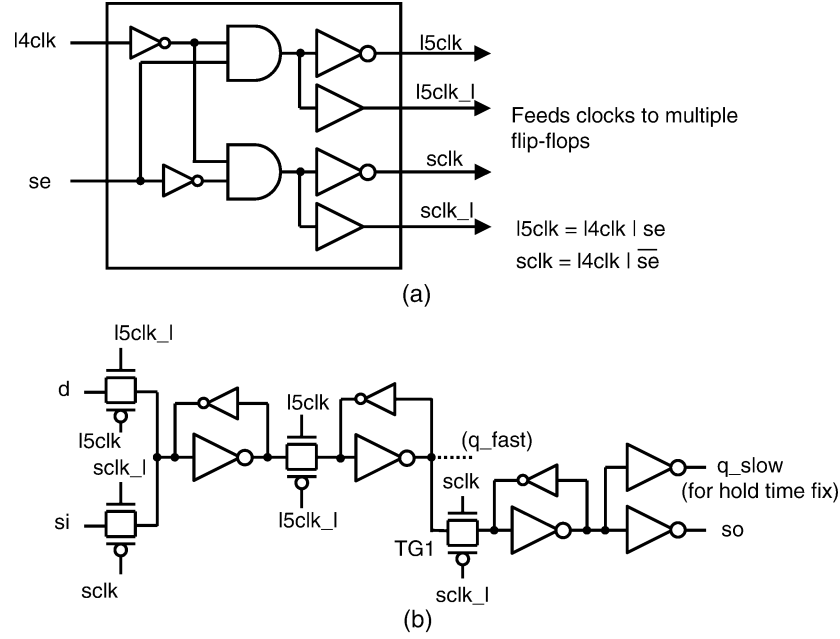


Fig. 13. Scan flip-flop with delayed output. (a) Clock header. (b) Base flip-flop.

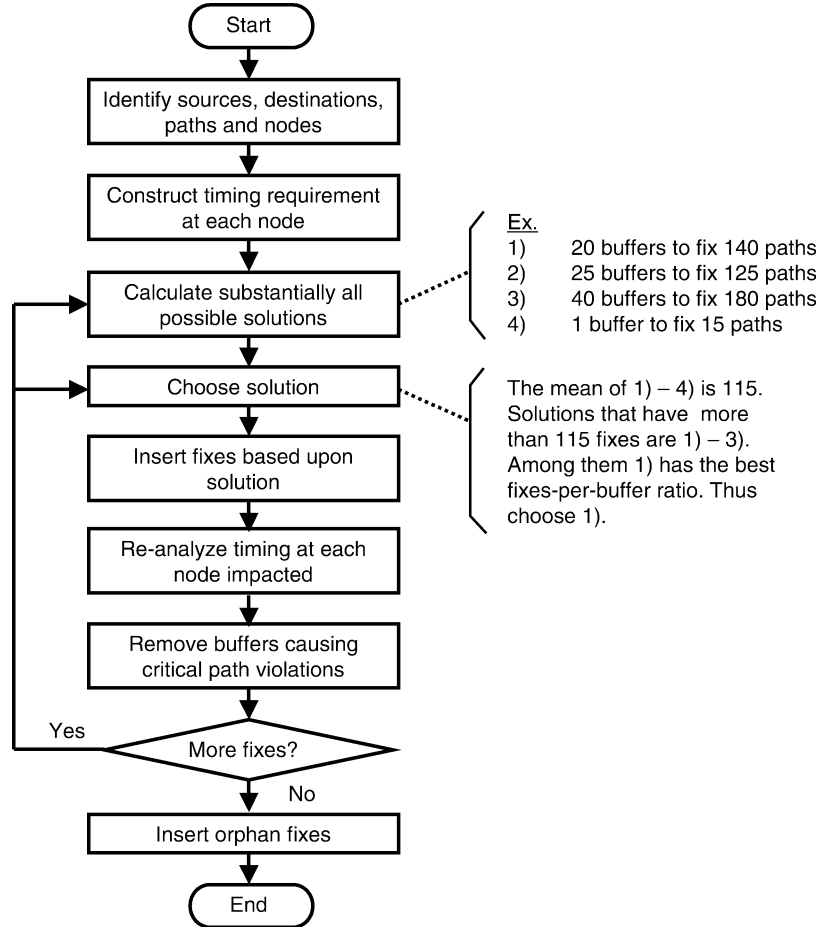


Fig. 14. Chart of hold-time fix flow.

E. Electromigration

Electromigration (EM) was another concern for this technology. The current density design rules were defined

for a 10-year EM lifetime. The feasibility study showed that, while the Cu metal layers scaled well from the previous technology, the new Cu vias did not scale well with respect to EM immunity. Both dc current and ac current in the signal

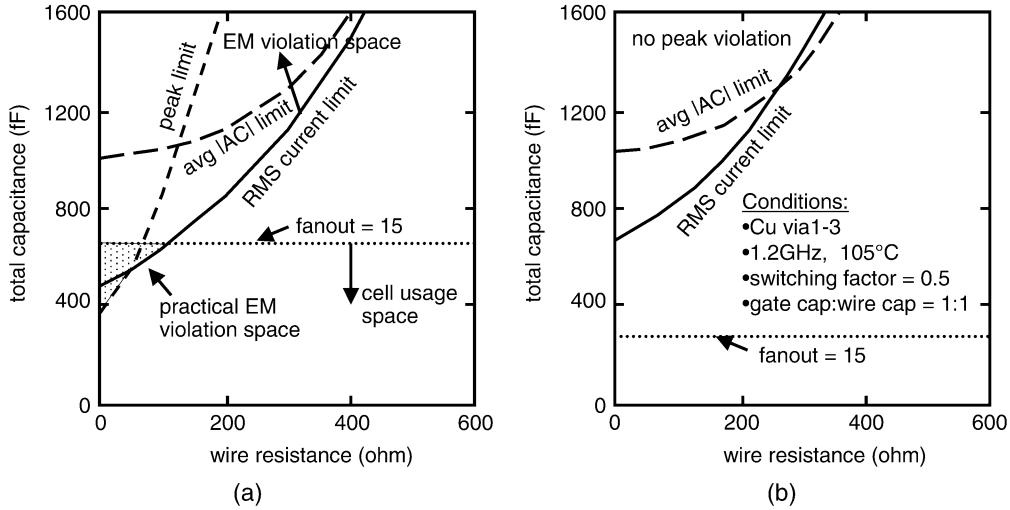


Fig. 15. Signal EM violation space versus cell usage space, showing that $32\times$ driver can cause an EM violation for rms and peak current limit in its usage space (fanout < 15), while $16\times$ driver does not cause the violation. (a) $32\times$ driver ($W_p/W_n = 26.9\mu\text{m}/13.4\mu\text{m}$), (b) $16\times$ driver ($W_p/W_n = 13.4\mu\text{m}/6.7\mu\text{m}$).

interconnect were of concern. Further analysis of EM constraints for average current, root-mean-square (rms) current and peak current identified certain combinations of driver size and RC load conditions that could cause signal EM violations (Fig. 15). Therefore, additional vias were implemented in the library cells with high drive strength without affecting the footprint. To reduce the design cycle and minimize the manual fix of EM violations on control blocks, EM analysis results were fed back to the place-and-route flow, which inserted more vias and widened the metals for the selected signals without creating DRC violations. The top-level router also was configured to place double vias for long interconnects and heavily loaded signals exceeding certain load thresholds.

IV. L2 CACHE IMPLEMENTATION

The increasing size of on-chip L2 caches results in longer latencies, as higher wire delay becomes a dominant factor in array access. In addition, soft error protection is increasingly essential for reliability due to the larger number of memory cells. Therefore, minimizing latency with ECC becomes a critical design factor for improving processor performance in larger on-chip L2 caches.

This chip uses two newly designed L2 caches. Each subsystem is a unified four way set-associative 512-kB L2 with integrated control and error correction logic. The cache line size is 64 bytes with a 128-bit data I/O bus, making a total of 2048 cache lines per set. The data is protected with 64 ECC bits for each line (8-bits per 64-bit datum) to support single-bit correction and double-bit detection. The 8-ECC per 64-bit combination gives the optimum tradeoff for this design between performance and overhead, including area and timing. The tag contains two bits for parity detection to flag soft errors as well as three modified-owned-exclusive-shared-invalid (MOESI) protocol bits to support cache coherency. Two random bits are stored for each cache line to support pseudo-random line replacement.

The L2 cache pipeline and floorplan are shown in Figs. 16 and 17. In cycle 1, addresses are dispatched from the core into

each array. A full cycle is necessary to send the addresses to the far end of the data arrays after multiplexing the addresses in the L2 control unit. In the next cycle, the data and tag arrays start the access simultaneously. In cycle 3, the address information from the tag array generates a way select (waysel) signal before the data arrives at the waysel datapath. This cycle ends by registering a selected set of data in a datapath block located in the center of the subsystem where the delay from all the data arrays is balanced. Cycles 2 and 3 are designed as multicycle paths for speed and power, eliminating the pipeline flops between the cycles, minimizing flop and clock skew overhead. In the last cycle, ECC syndrome generation and correction is performed in the datapath located on the channel where data is routed back to the core. The final data set is registered in the interface block located at the top of the L2 subsystem. This design achieved a low four-cycle latency from the L2 cache to the core including error correction.

The control logic placed in the center of the subsystem adjacent to the tag arrays handles several functions, including partial stores for read-modify-write of 8-bit per 64-bit ECC algorithm, way locking and delayed write buffers. The datapath logic blocks, including way-selection, ECC and interfaces to the core and system bus, are optimally placed along with the data-flow in the center of the subsystem and in the channel region between the SRAM arrays. This approach of integrating the datapath and control logic closely with the SRAM arrays, and place-and-routing them based on their timing criticality, reduces the required number of pipeline stages for communication between the arrays and the logic blocks by minimizing the impact of long wire delays. This solution contrasts with a conventional design approach, where the logic is partitioned and placed separately from the SRAM arrays. In addition, as the logic blocks are located close to the SRAMs, optimizing interconnect overhead, power consumption is reduced on the information exchange between SRAM arrays and logic blocks.

The SRAM arrays are designed to support basic read and write operations only, simplifying the array design itself and reducing test complexities. The comparator and ECC logic are

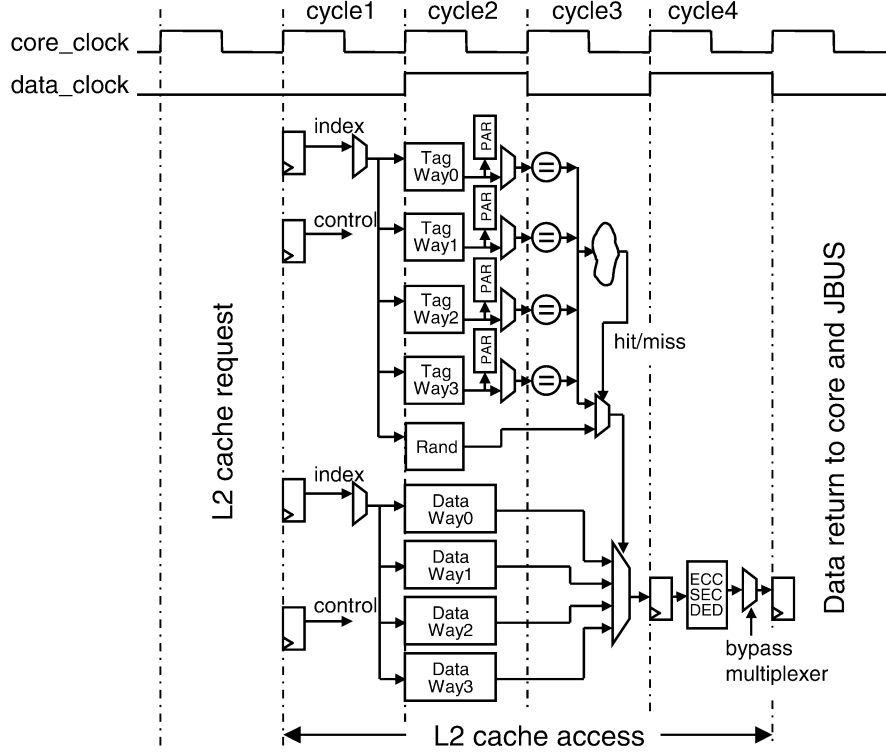


Fig. 16. L2 cache pipeline diagram.




Fig. 17. L2 cache subsystem floorplan.

realized in the datapath blocks with automated design flows instead of using a custom implementation inside the SRAMs. This approach not only significantly reduced the initial design time but also facilitated subsequent iterations of the design for timing closure.

V. INTEGRATION

Roughly 14% of the control and datapath blocks in the UltraSPARC core area were re-synthesized, placed and routed, while other blocks were migrated to the new technology with engineering change orders (ECOs) to achieve a faster floorplan and timing closure. The core level routing was redone to deal with the deep submicron interconnect issues. Tradeoffs for area, speed, coupling noise, and EM were optimized by the use of various wire classes, defined in terms of width, space, and shielding (half-side or both-side shielded). Repeater insertion was also performed for speed, coupling noise, EM and signal slew rate. At the core level, 2400 new repeaters were inserted, while 12 900 repeaters were implemented at the chip level.

An enhanced version of the in-house CAD tool described in [12] was applied to this design to exhaustively analyze the coupling noise in interconnects. As the tool first performs a worst-case analysis, assuming all aggressors switch simultaneously, the reported results are inherently conservative. Accounting for the logic and timing constraints of the victim and aggressor nets effectively filtered out the false violations.

To improve critical paths, low- V_t transistors are applied for 6% of the total transistor width at the core, 3% total at the chip level. Footprint compatible low- V_t library cells, coupled with an automated cell replacement flow, made this process more efficient. Under nominal operating conditions, the increase in leakage power due to the low- V_t transistors is less than 200 mW.

VI. SILICON RESULTS AND CONCLUSION

A highly integrated 64-bit dual-core microprocessor optimized for low-cost dense servers was successfully created.




Fig. 18. Shmoo plot at 85°C.

The target frequency of 1.2 GHz at 1.3 V, 85°C was achieved with comfortable margin (Fig. 18), dissipating 23 W with typical application workloads. High performance-to-power ratio is realized by reusing the power and area efficient UltraSPARC II core. The core circuits were successfully transported from a 0.25- μ m process to a 0.13- μ m process, with high circuit margins to cope with new deep submicron design issues. A low-latency and high-capacity L2 cache was newly designed. The four-cycle core-to-L2 latency including ECC is short for 64-bit processors. A short design cycle was achieved by leveraging existing designs and developing highly effective design methodologies and flows. The chip was taped out slightly over one year after the project was fully staffed.

ACKNOWLEDGMENT

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 R. Heald, J. Kaku, and H. McGhan for their technical advice, and A. Kowalczyk, R. Raman, A. Strong, F. DeSantis, H. Ward, and J. England for supporting the project.

REFERENCES

- [1] M. Horowitz and W. Dally, "How scaling will change processor architecture," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2004, pp. 132–133.
- [2] S. Kapil, "Gemini: A power-efficient chip multithreaded UltraSPARC processor," presented at the 15th Hot Chips Symp., Stanford, CA, Aug. 2003.
- [3] T. Takayanagi *et al.*, "A dual-core 64-bit UltraSPARC microprocessor for dense server applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2004, pp. 58–59.
- [4] L. A. Lev *et al.*, "A 64-bit microprocessor with multimedia support," *IEEE J. Solid State Circuits*, vol. 30, no. 11, pp. 1227–1238, Nov. 1995.

- [5] D. Greenhill *et al.*, "A 330 MHz 4-way superscalar microprocessor," *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 166–167, Feb. 1997.
- [6] G. K. Konstantinidis *et al.*, "Implementation of a third-generation 1.1-GHz 64-bit microprocessor," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1461–1469, Nov. 2002.
- [7] E. Anderson, "A 64-entry 167 MHz fully-associative TLB for a RISC microprocessor," *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 360–361, Feb. 1996.
- [8] E. Seevink *et al.*, "Current-mode techniques for high-speed VLSI circuits with application to current sense amplifier for CMOS SRAM's," *IEEE J. Solid-State Circuits*, vol. 26, no. 4, pp. 525–536, Apr. 1991.
- [9] N. Shibata, "Current sense amplifiers for low-voltage memories," in *IEICE Trans. Electron.*, vol. E79-C, Aug. 1996, pp. 1120–1130.
- [10] V. Reddy *et al.*, "Impact of negative bias temperature instability of digital circuit reliability," in *Proc. Int. Reliability Physics Symp.*, 2002, pp. 248–254.
- [11] N. Kimizuka *et al.*, "The impact of bias temperature instability for direct-tunneling ultra-thin gate oxide on MOSTET scaling," in *Symp. VLSI Technology Dig. Tech. Papers*, Jun. 1999, pp. 73–74.
- [12] K. Aingaran *et al.*, "Coupling noise analysis for VLSI and ULSI circuits," in *ISQED 2000 Proc.*, Mar. 2000, pp. 485–489.




Toshinari Takayanagi (M'94) was born in Aichi, Japan, in 1962. He received the B.S. and M.S. degrees in electrical engineering from the University of Tokyo, Tokyo, Japan, in 1985 and 1987, respectively.

In 1987, he joined Toshiba Semiconductor Device Engineering Laboratory, Kawasaki, Japan, where he was engaged in high-speed circuit design of memories and arithmetic units, low-power circuit design, CMOS RF circuit design, development of microprocessor, media-processor and wireless communication LSIs, and design methodologies.

His work includes off-chip/on-chip cache SRAMs, structured array, R8000 4-way superscalar microprocessor, PS2 Emotion Engine processor, MPACT media processor, an MPEG4 videophone LSI, and a Bluetooth LSI. In 2001, he joined the Processor Development Group, Sun Microsystems, Sunnyvale, CA, where he has been engaged in high-performance and high-throughput SPARC microprocessor designs, including dual-core UltraSPARC Gemini and next-generation CMT (Chip MultiThreading) processor Niagara. His current research interests include co-optimization of system design, circuit design and device/process technologies. He has authored or coauthored 24 conference and journal papers in the VLSI design field and holds 16 patents.

Mr. Takayanagi serves as a technical committee member of the International Conference on IC Design and Technology (ICIDT).



Jinuk Luke Shin (M'98) received the B.S. degree from the University of Louisville, KY, with highest honors, and the M.S. degree from the University of Texas, Austin, both in electrical engineering.

He was with Motorola, Austin, from 1995 to 1997, where he was involved in the design of embedded flash memories for digital signal processors. From 1997 to 2000, he was with Hitachi Semiconductor America, San Jose, CA, where he developed the cache array for a SuperH processor. He joined Sun Microsystems Inc., Sunnyvale, CA, in 2000. He was


the technical lead and manager for the L2 cache and TLB designs of the dual core SPARC processor. He is currently a Senior Design Manager and lead for the level 1 and 2 caches and the analog blocks for Sun's next-generation multithread microprocessor. He is a co-author of 15 technical papers in design, device, and test areas, and holds four U.S. patents with several pending.

Mr. Shin is a member of Eta Kappa Nu and Tau Beta Pi.




Bruce Petrick received the B.S. and M.S. degrees in electrical engineering from Montana State University, Bozeman, in 1972 and 1974, respectively.

He joined Sun Microsystems in 1994 and worked as a Design Engineer on the UltraSPARC II, later working on the UltraSPARC III and most recently the Gemini chip. He was responsible for the design of the integrated level 2 caches and the logic changes in the CPU core. Currently, he is working on a new multi-thread CPU design.



Nathan Moon received the M.S. degree in electrical engineering from Texas A&M University, College Station, in 1986.

He has been with Sun Microsystems since 2002, working on the custom circuits, SRAMs, CAMs, and clock. He is currently designing TCAM for the next-generation UltraSPARC microprocessor. Prior to joining Sun, he was mainly involved in circuit design, testing, and productization for nonvolatile memories at Texas Instruments and Motorola.



Jeffrey Y. Su received the M.S. degree in electrical engineering from University of California, Irvine, in 1989.

He joined Sun Microsystems, Austin, TX, in 2001, where he managed a circuit design team responsible for SRAM, clock, and custom circuit design for UltraSPARC projects. Currently, he is leading clock design on the next-generation UltraSPARC microprocessor. Prior to joining Sun, he was involved with flash memory circuit design at Motorola. His current interests include memory compiler, clock

distribution, clock skew optimization, and DFT.



Dina Bistry received the B.S. degree in computer science from the Technion-Israel Institute of Technology in 1979.

She joined Sun Microsystems, Sunnyvale, CA, in 1998. Prior to that, she worked over ten years for Intel in Santa Clara, CA, and Israel. She specializes in various aspects of performance verification, first as a CAD tool developer, later as an application engineer and methodology developer. She was responsible for the performance verification of Gemini.



Howard Levy received the B.S.E.E. degree from The Ohio State University, Columbus, in 1994.

After working for IBM for seven years and Intrinsity for one year in the area of circuit and memory design, he joined Sun Microsystems in 2000. Since working with Sun, he has worked on UltraSPARC II, Niagara I and Niagara II processors in the area of memory and CAM design. He holds five U.S. patents in the area of high-speed circuit design.



Umesh Nair received the M.S. degree in electrical engineering from Texas A&M University, College Station, in 1998.

He has been with Sun Microsystems since 1998, engaged in timing verification and physical design activities. He is currently working on the next-generation CMT (Chip MultiThreading) processors. His research interests are focused on the physical design aspects of high-speed processors.




Ha Pham received the B.S. degree in electrical engineering and computer science from the University of California at Berkeley in 1996.

Since 1998, he has been working at Sun Microsystems Inc., Palo Alto, CA, completing a number of UltraSPARC microprocessors. He is currently a Staff Engineer working on a new microprocessor based on throughput computing architecture. Prior to joining Sun, he worked at Intel Corporation and Silicon Storage Technology Inc.



Mandeep Singh received the B.E. degree in electronics and electrical communication from Panjab University, India, and the M.S. degree in computer engineering from University of Cincinnati, Cincinnati, OH, in 1996 and 1999, respectively.

Since 1999, he has been with Sun Microsystems working on the physical design including noise, electromigration, IR drop analysis and full chip integration for several UltraSPARC processors. Currently, he is the Integration Lead for the next-generation multithreading SPARC processor.



Jinseung Son received the B.S. degree in electronic engineering from Seoul National University, Korea, in 1991, and the M.S. degree in electronic engineering from Pohang University of Science and Technology, Korea, in 1993.

In 1993, he joined Hynix Electronics, Korea, where he was involved in the development projects of prototype 256 M SDRAM and 16 M RAMBUS DRAM. From 1998 to 2000, he worked for Mosel Vitelec, San Jose, CA, where he designed 16 M and 64 M SDRAMs. Since 2000, he has been with

Sun Microsystems, where he has been involved in the Cache designs of L2 Data Array and L1 Instruction Data Array for Sun's SPARC CPUs. His current research interests include memory embedment techniques for high performance and density.



Vikas Mathur received the B.Eng. degree in electronic engineering from Panjab University, India, in 1996, and the M.S. degree in electrical engineering from University of Kentucky, Lexington, in 1998.

He joined ASIC design group of Apple Computers, Cupertino, CA, in 1998, where he worked on design verification and hardware-software co-simulation. In 2000, he joined the UltraSPARC microprocessor group at Sun Microsystems Inc., where he has been involved in back end and SRAM design.



Ana Sonia Leon (S'83–M'88) received the B.S. degree in electrical engineering from the Polytechnic University of Ecuador (ESPOL) in 1987, and the M.S. degree in electrical and computer engineering from the University of Southern California, Los Angeles, in 1988.

From 1989 to 1996, she was a Lead Circuit Designer with Motorola, Inc., Austin, TX, working in the design of DSP processors with focus on digital circuits, clock and power distribution, phase-lock-loops and design methodology. In 1996 she joined Chromatic Research, Sunnyvale, CA, as a Member of the Technical Staff, where she was engaged in the development of several generations of MPACT multimedia processors, in the areas of circuit design, integration, and design methodology, and led the implementation and tapeout of several chip designs. Since 1998 she has been with Sun Microsystems, Inc., Sunnyvale, CA, where she has led and managed the backend implementation for several processors including the UltraSPARC IIe Hummingbird and the dual-core UltraSPARC Gemini. Her engagement includes SRAMs and megacells designs, global circuits, integration, and design methodology. She is currently a senior manager responsible for the backend design and implementation of the next generation of CMT (Chip MultiThreading) processor Niagara. Her technical interests include high-speed interconnect and clocking, power distribution, power, and leakage reduction, and high-speed design methodology. She holds seven issued patents in circuit design and is the co-author of several technical microprocessor papers.

Ms. Leon is a member of Tau Beta Pi and Eta Kappa Nu.