

# A 480-MHz RISC Microprocessor in a 0.12- $\mu\text{m}$ $L_{\text{eff}}$ CMOS Technology with Copper Interconnects

Chekib Akrou, John Bialas, Miles Canada, Duane Cawthron, James Corr, Bijan Davari, *Senior Member, IEEE*, Robert Floyd, Stephen Geissler, Ronald Goldblatt, Robert Houle, Paul Kartschoke, Diane Kramer, Peter McCormick, *Member, IEEE*, Norman Rohrer, Gerard Salem, Ronald Schulz, Lisa Su, *Member, IEEE*, and Linda Whitney

**Abstract**—This paper describes the performance improvements of a reduced instruction set computer (RISC) microprocessor that has migrated from a 2.5-V technology to a 1.8-V technology. The 1.8-V technology implements copper interconnects and low  $V_t$  field-effect transistors in speed-critical paths and has an  $L_{\text{eff}}$  of 0.12  $\mu\text{m}$ . Global clock latency and skew are improved by using copper wires, and early mode timings are improved by reducing clock skew and adding buffers. These enhancements, along with an environment of 2.0 V, 85°C, and with a fast process, produced a 480-MHz RISC microprocessor.

**Index Terms**—CMOS, copper, low threshold, microprocessor, PowerPC, reduced instruction set computer (RISC).

## I. INTRODUCTION

A 32-BIT 480-MHz PowerPC<sup>1</sup> reduced instruction set computer (RISC) microprocessor has been migrated into an advanced 0.2- $\mu\text{m}$  CMOS technology with copper for its interconnects and multithreshold transistors. The technology features help to increase the microprocessor's internal clock frequency up to 480 MHz at 2.0 V and 85°C, and at the fast end of the process distribution. When operating at room temperature, the clock frequency increases to over 500 MHz.

## II. PROCESSOR ARCHITECTURE

The microprocessor's architecture can dispatch two instructions per cycle by using two 32-KB L1 caches, one for data and one for instructions. An integrated L2 cache controller can work with L2 cache sizes of 256 KB, 512 KB, or 1 MB, and the I/O's interface with the external bus using an industry standard of 3.3 V. The 60 $\times$  bus has ratios of 1 $\times$  as well as 2 $\times$  to 8 $\times$  with 0.5 $\times$  increments. The L2 cache interface contains ratios of L2 bus frequency to internal clock frequency of 1:1, 2:3, 1:2, 2:5, and 1:3. The microprocessor has been successfully implemented in a 2.5-V CMOS technology [1] and has migrated to a 1.8-V CMOS technology (Fig. 1). The 1.8-V microprocessor has estimated

SPECint95 and SPECfloat95 ratings of 20 and 12, respectively, at 480 MHz with a 1-MB L2 cache running at 240 MHz and a 60 $\times$  bus frequency of 96 MHz. The 60 $\times$  bus is the high-speed 32- and 64-bit processor bus interface used throughout the PowerPC family of processors.

## III. TECHNOLOGY

The 2.5-V CMOS technology, with a field-effect transistor (FET)  $L_{\text{eff}}$  of 0.18  $\mu\text{m}$ , has five interconnect levels made of aluminum, a tungsten local interconnect, and tungsten vies. The local interconnect contacts both diffusion and polysilicon by overlap. The contacted wiring pitch for the third level of metal is 1.26  $\mu\text{m}$ . The metal thicknesses are 0.54  $\mu\text{m}$  for M1, 0.73  $\mu\text{m}$  for M2, M3, and M4, and 2.07  $\mu\text{m}$  for M5. The nominal internal clock frequency is 221 MHz; but the frequency increases to greater than 275 MHz when the  $L_{\text{eff}}$  line is tailored shorter, the voltage is increased, and the maximum operating temperature is reduced. The I/O's were designed to interface from the internal voltage of 2.5 V to the 3.3-V external bus. The 6.4 million transistors fit into a 67 mm<sup>2</sup> chip area.

The 1.8-V CMOS technology has a nominal NFET  $L_{\text{eff}}$  of 0.12  $\mu\text{m}$ . This technology is a linear shrink from the 2.5-V technology with some changes to the ground rules for the local interconnect and the Si/local interconnect spacing. In this technology, the microprocessor has six levels of metal (Fig. 2). Metal-level thicknesses are 0.40  $\mu\text{m}$  for M1 and M2, 0.55  $\mu\text{m}$  for M3 and M4, and 1.20  $\mu\text{m}$  for M5 and M6. Beginning with M1, all levels of metal and vias are fabricated with copper. Again, a tungsten local interconnect with a tungsten via is used to increase the overall density. The copper wires, manufactured with a damascene process, are deposited via electroplating, which provides lower costs and better connections to the vias than the tungsten via and metal etch-back process used for aluminum. A single damascene process is used for M1 and a dual damascene process is used for the remaining metal levels. The dual damascene process produces the via connections and the metal level, in copper, at the same time. The third-metal level has a reduced contacted pitch of 0.81  $\mu\text{m}$ . This microprocessor's SRAM cell, used in several custom circuits, is a migrated design with new ground-rule waivers measuring 7.6  $\mu\text{m}^2$ . The smallest SRAM cell available in the 1.8-V technology, which measures 6.84  $\mu\text{m}^2$ , was not selected in order to maintain the M1 word line and not adversely affect the development schedule. The technology offers FET's with two

Manuscript received July 10, 1998; revised September 22, 1998.

C. Akrou, D. Cawthron, D. Kramer, and L. Whitney are with IBM Microelectronics Division, Austin, TX 78758 USA.

J. Bialas, M. Canada, J. Corr, R. Floyd, S. Geissler, R. Houle, P. Kartschoke, P. McCormick, N. Rohrer, and G. Salem are with IBM Microelectronics Division, Essex Junction, VT 05452 USA (e-mail: pkartsch@btv.ibm.com).

B. Davari, R. Goldblatt, R. Schulz, and L. Su are with IBM Microelectronics, Hopewell Junction, NY 12533 USA.

Publisher Item Identifier S 0018-9200(98)07042-5.

<sup>1</sup>PowerPC is a trademark of IBM Corporation.

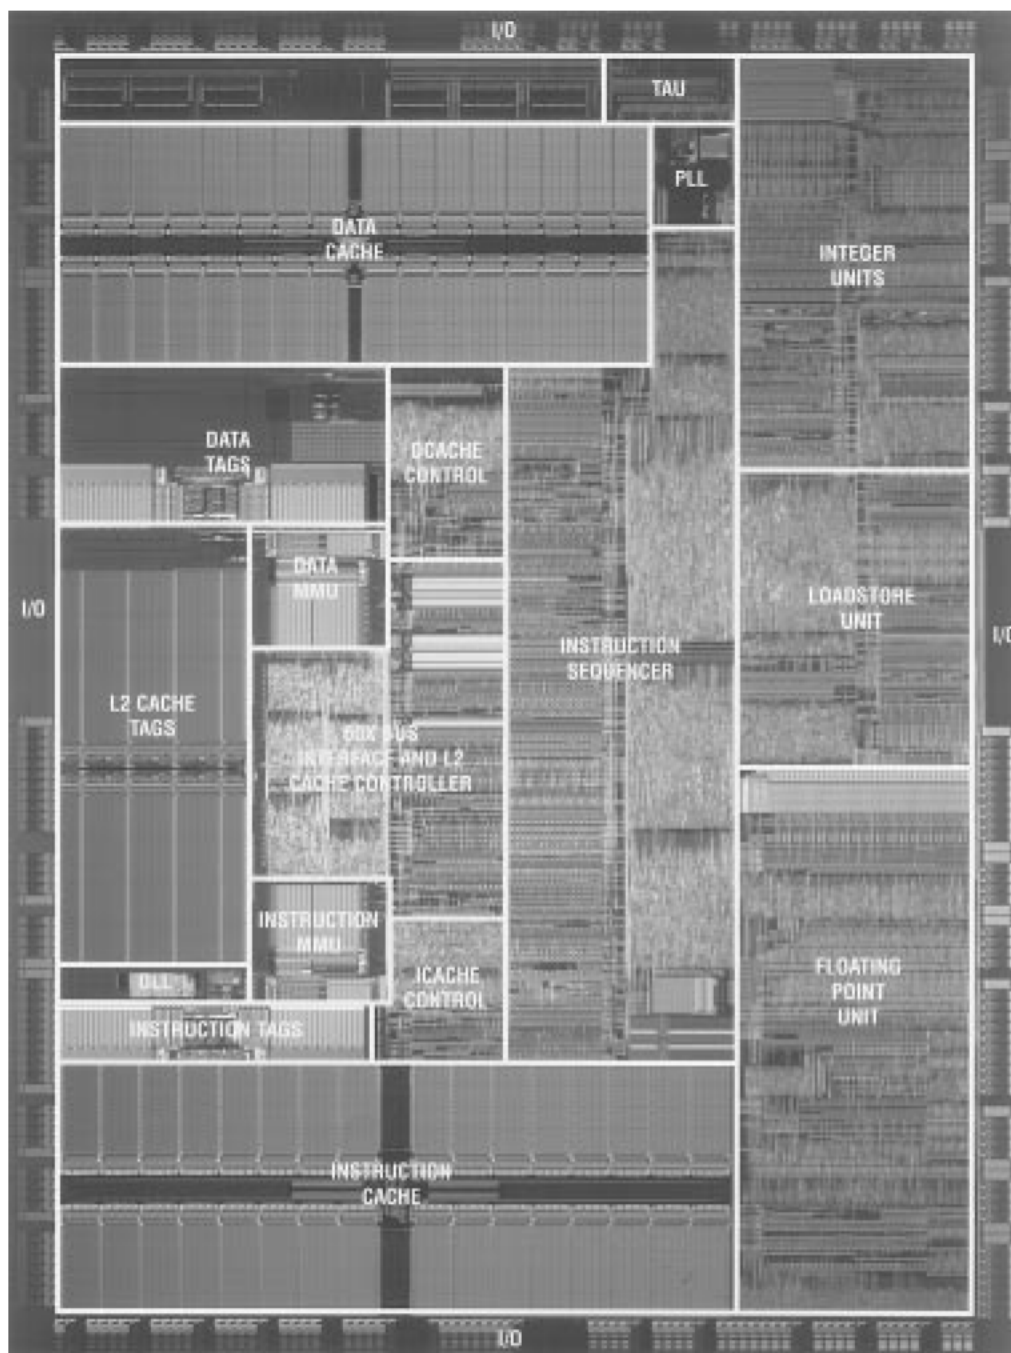


Fig. 1. A 480-MHz RISC microprocessor with copper interconnect.

distinct threshold voltages ( $V_t$ ) for both NFET's and PFET's. At long channel lengths, the low threshold voltage NFET has a  $V_t$  that is 100 mV lower than the nonlow  $V_t$  device of 310 mV. The low  $V_t$  devices are provided to increase performance. The microprocessor's present chip size is 40 mm<sup>2</sup>. Table I lists key microprocessor and technology parameters for both the 2.5- and 1.8-V technologies.

#### IV. COPPER INTERCONNECT

Copper interconnects provide a 40% decrease in RC delay over the traditional aluminum wires due to the reduced resistivity of copper [2]. Reduced copper resistivity allows for a

tradeoff of the resistance or capacitance by thinning the wire. The sheet resistance of the previous interconnect technology is maintained because a thinner wire improves manufacturability and yield and because the aspect ratio of the wire is not as extreme. Reliability of these thinner wires is not compromised because copper has less electromigration than aluminum. For the 1.8-V technology, the RC reduction of copper is mainly caused by the reduced capacitance; the sheet resistance for a given interconnect remains nearly constant. The total RC reduction is 23%. Because of lower capacitance on the wires, intralayer coupling capacitance between the third-metal-level wires is reduced by 27%. Fig. 3 shows the late-mode timing delay of all the timing paths when the microprocessor is

TABLE I  
TECHNOLOGY FEATURES

Property of Microprocessor	2.5V Technology (Aluminum)	1.8V Technology (Copper)
Voltage	$2.5 \pm 0.2V$	$1.8 \pm 0.15V$
I/O Voltage	$3.3 \pm 0.3V$	$3.3 \pm 0.3V$
L2 Cache I/O Voltage	3.3V	3.3V or 1.8V
$L_{eff}$ for NFET/PFET	$0.18\mu m/0.18\mu m$	$0.12\mu m/0.15\mu m$
$T_{ox}$	5.0nm	3.5nm
Contacted Local Interconnect Pitch	$0.84\mu m$	$0.55\mu m$
Contacted M1 Pitch	$0.98\mu m$	$0.63\mu m$
Contacted M2, M3, M4 Pitch	$1.26\mu m$	$0.81\mu m$
Contacted M5, M6 Pitch	$2.52\mu m$	$1.62\mu m$
M1 Thickness	$0.54\mu m$	$0.40\mu m$
M2 Thickness	$0.73\mu m$	$0.40\mu m$
M3 Thickness	$0.73\mu m$	$0.55\mu m$
M4 Thickness	$0.73\mu m$	$0.55\mu m$
M5 Thickness	$2.07\mu m$	$1.2\mu m$
M6 Thickness		$1.2\mu m$
SRAM Cell Area	$18.4\mu m^2$	$7.6\mu m^2$
Chip Size	$67mm^2$	$40mm^2$
Power	5.4W @ 300MHz	3.0W @ 400MHz




Fig. 2. Scanning electron microscope of six copper metal levels and local interconnect.

wired with aluminum. Fig. 4 presents the same data with the metal levels fabricated in copper; the path delay near 2.0 ns is reduced because the RC delay percentage on the critical paths is decreased. Path delays less than 1.0 ns did not change significantly because the paths contained little RC delay. The top 5% of the paths have RC delay ranging from 5–25% of the total path delay, where the speed advantage




Fig. 3. Histogram of path delays with Al interconnects.

from copper is noticeable. The top 5% of the paths typically contain 18–23 gates/cycle. With a higher percentage of the critical path delay dominated by intrinsic gate delay, the frequency of the microprocessor more easily increases at the fast end of the process distribution, with either higher voltage or lower temperature. Some of the late-mode critical timing paths, which still contain a large amount of RC, are tripled in width to decrease the RC delay further. Decreased cycle time is accomplished without adding a pipe stage, so a high frequency can be reached while still preserving the number of instructions executed per cycle. A similar timing for early mode was performed to compare clock paths to shortest delay




Fig. 4. Histogram of path delays with Cu interconnects.

data paths. If a latching element has data inputs that transition too fast or a clock that transitions too slow, the latch can be corrupted. Because copper interconnects reduce RC delay and can produce early mode paths, buffers are added to slow down and eliminate these paths.

#### V. LOW $V_t$ FET IMPLEMENTATION

The microprocessor's internal clock frequency is further improved with low  $V_t$  FET's. With our timing and floor-planning methodology, low  $V_t$  transistors are inserted manually into custom macros to improve delay through a custom circuit. The remaining logic on the chip is implemented with 4-bit nibbles of random logic gates (extended custom standard-cell library) and individual 1-bit standard-cell books. A duplicate low  $V_t$  library has been developed for these standard cells. To allow nonlow  $V_t$  and low  $V_t$  books to be interchangeable, the layout of the low  $V_t$  books is identical to that of the nonlow  $V_t$  books in size and wiring. After timing the chip at 500 MHz, any part of a late-mode timing path that contains standard-cell books and does not meet the timing requirements is converted to a low  $V_t$  book and the chip is retimed. Because low  $V_t$  books are used only for the critical paths, the number of transistors with the reduced threshold voltage is limited. For the whole chip, 12.3% of the 4-bit standard-cell books and 5.0% of the 1-bit standard-cell books were converted to low  $V_t$  books. The  $I_{dsat}$  for the low  $V_t$  FET's is 10% higher than for the standard FET's. The low  $V_t$  FET leakage is 8–10 $\times$  higher than standard FET's. Therefore, to balance power versus speed, only 4.2% of the total transistors on the microprocessor are converted to low  $V_t$ , which accounts for a performance increase of 6.5%. To minimize risk and maintain noise margins, this design does not implement low  $V_t$  FET's within dynamic circuits.

#### VI. CLOCK DISTRIBUTION AND PHASE-LOCKED LOOP DESIGN

Increasing the processor speed with a technology remap required a proportional improvement in the clock design and distribution because clock skew subtracts from the processor cycle time. A hierarchical clock-distribution methodology included variable wire widths to minimize delay, automated clock load balancing, and local clock buffer clustering. The clock distribution network contains four stages: a large global clock buffer, a clock tree, 500 local clock splitters, and local




Fig. 5. Histogram of the clock tree latency with Al and Cu interconnects.

clock buffers. The local clock splitter converts a global single-phase clock into complementary unbuffered C1 and C2 clocks. The local clock buffers use the split outputs to produce a buffered master latch clock C1 and a buffered slave latch clock C2.

The clock tree is designed in a pseudogrid-spine network. Thick metal levels M6 and M5 are used for the primary global clock wiring. Four M5 spines run parallel for 3 mm, and M6 is used to connect the spines to form a grid network. Wiring levels M4 through M2 are used to connect clusters of clock splitters to the spines, and total wire lengths are restricted to 500  $\mu$ m. If the 2.5-V-technology microprocessor implements six levels of metal, the improved RC delay from copper interconnect reduces clock latency from 170 to 85 ps and improves the clock tree skew from 55 to 35 ps, as shown in Fig. 5. This improvement in clock skew helps reduce the number of early mode fast paths. The copper interconnect allows the use of narrow metal lines, which reduces inductance effects; an aluminum interconnect could produce similar RC skew by using wider lines, but the increased inductance would affect skew significantly.

Because additional clock skew directly impacts the process cycle time, the local clock splitter and local clock buffer must deliver accurate clock signals to the latches. Therefore, the number of local clock buffers was increased from eight to 24, with each buffer driving a smaller restricted load range. The restricted buffer load capacitance range reduces clock buffer skew from 100 to 20 ps. Software tools allow automated repowering based on wire and latch capacitances. The low RC interconnect and smaller clock buffer load capacitance range improves the distribution of a falling C1 edge at the latch boundaries, as shown in Fig. 6, which compares the 2.5-V technology to the 1.8-V technology.

Critical components must be adjusted when a phase-locked loop (PLL) is remapped from a 2.5-V to a 1.8-V technology. The PLL [3] was originally designed for a 2.5-V CMOS technology. The critical components examined during a technology remap are the phase detector gain, voltage reference center current, charge-pump current sources, feed-forward current sources, and voltage-controlled oscillator (VCO) gain. Statistical analysis is performed for various parameter changes such as channel length, threshold voltage, temperature, power supply, and process mistracking deviations. Open-loop and






Fig. 6. Clock skew of all falling C1.

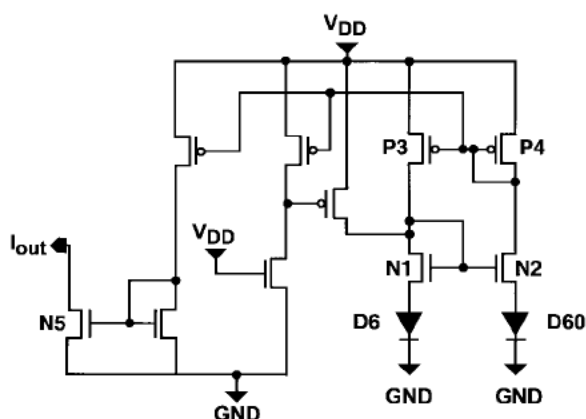


Fig. 7. Voltage reference circuit.

closed-loop circuit simulations as well as closed-loop behavioral time-domain models are used to optimize the PLL for damping coefficient and to improve stability of the VCO pole, phase error, and jitter.

Analog circuit operation becomes difficult when the power supply is scaled because MOS device threshold voltage and dc operating points are not scaling proportionally. The voltage reference circuit consists of a pair of ratioed diodes (D6 and D60), a pair of ratioed NMOS transistors (N1 and N2), and a pair of load PMOS transistors (P3 and P4), as shown in Fig. 7. As the technology is scaled, the diode forward bias remains constant, the MOS device dc bias condition ( $V_{gs} - V_t$ ) scales by  $0.85\times$ , and the power supply scales by  $0.72\times$ . Thus, as the power supply is lowered, the voltage-reference-transistor operating points move from saturation into the linear region. When the transistors are in the linear region of their operation, the charge pump stops functioning as an infinite current source and operates with poor stability over process, temperature, and voltage. The voltage-reference-transistor operation in the saturation region is guaranteed by converting NMOS transistors (N1 and N2) from standard threshold devices to zero  $V_t$  devices and PMOS transistors (P3 and P4) from standard  $V_t$  to low  $V_t$  devices. The improvement in the voltage reference circuit is measured through transistor N5 (Fig. 8).

## VII. TIMING AND PERFORMANCE

The 1.8-V technology improves the frequency of the processor by 27%, while the logic, circuit, and floor-planning




Fig. 8. Voltage reference current source N5.





Fig. 9. Typical path delays.

Fig. 10.  $F_{max}$  versus voltage for critical path delays.

enhancements contribute a 17% improvement. In addition, the copper wires for this design result in a 12% improvement in path delay, while low  $V_t$  devices add another 6.5% to the maximum frequency. Total improvement is 77%. One of the critical path delays is shown in Fig. 9. The timing delay is segmented into custom gate delay, standard-cell gate delay, and interconnect delay for the 2.5-V technology, 1.8-V technology without copper interconnects or low  $V_t$  FET's, 1.8-V technology with copper interconnects and low  $V_t$  devices, and 1.8-V technology with circuit enhancements.

A timing analysis was performed at 480 MHz, 2.0 V, 85°C, and all process parameters were centered at 1-sigma fast—except for  $L_{eff}$ , which was centered at 3-sigma fast—to represent a chip design line-tailored to the fast corner. The fast corner of the process distribution also shows more RC effects. Fig. 10 shows the late-mode timing for revisions 1.0 and 2.0 of the 1.8-V microprocessor and the timing of the




Fig. 11. Internal and L2 clock timings.

2.5-V microprocessor; the slope of the revision 2.0 data is 20 MHz/100 mV.

### VIII. SYSTEM AND L2 CACHE INTERFACE

A new option for the L2 bus interface with the L2 cache is featured on the 480-MHz microprocessor. The present 3.3-V interface is maintained with an improved I/O that allows the internal voltage to be reduced to 1.8 V without creating additional stress on the thin gate oxide in the drivers; this I/O also provides good 50- $\Omega$ m impedance matching to the external pins. Fig. 11 displays the internal clock to 60 $\times$  bus ratio of 6:1 and the L2 clock to an internal clock ratio of 1:2. When the L2 cache is fabricated in a 1.8-V technology, a second L2 bus I/O can interface between the 1.8-V internal supply and the 1.8-V external L2 bus by changing the vies between M2 and M3; this enables a faster L2 bus data transfer between the microprocessor and L2 cache. When using the 1.8-V receivers and drivers in a multichip module or multimodule card, the L2 cache and the microprocessor are expected to interface with a 1:1 or 2:3 ratio of L2 bus frequency to internal clock frequency.

The improved I/O incorporates several advances in mixed-voltage I/O interface techniques. The off-chip driver is designed to provide low voltage transistor transistor logic (LVTTL)-compliant off-chip interface levels with a total swing of 0–3.3 V, using a dual power supply of the 1.8-V core and a 3.3-V second I/O supply. By incorporating voltage-dividing techniques to protect individual transistors from overvoltage stresses, the I/O circuit design provides the capability of supplying a 3.3-V output level using conventional transistors in the 1.8-V technology. Key enhancements over prior designs include a “pulsed” voltage translation circuit that converts a 0–1.8 V signal swing into a 1.8–3.3 V signal swing for output device predrive. This pulsed circuit eliminates the problems of prior voltage-translation designs that exhibit both poor throughput performance and large dc power consumption. Also, an additional circuit provides a precisely controlled reference voltage to drive the final output transistors, which eliminates a sensitivity to the voltage separation of the core and I/O supply identified in prior designs and minimizes power consumption. This reference voltage circuit renders the design virtually immune to performance—and impedance-matching variations caused by the collapse of the difference between the core and I/O supplies because the core voltage is increased to improve core logic performance. The driver

is also source-terminated for impedance control over a wide process range, and the source-termination technique inherently provides a robust electrostatic discharge protection network. The on-chip receiver is also enhanced with a zero  $V_t$  input clamp device, giving overvoltage stress protection to sensitive receiver circuits and guaranteeing that the receiver is LVTTL-compliant.

### IX. CONCLUSIONS

The 480-MHz RISC microprocessor has successfully demonstrated improvements when migrating from a 2.5-V technology to a 1.8-V technology. Design improvements include additional buffers for early mode, circuit-design changes, global clock latency, and skew reduction. Technology improvements include a copper interconnect with a 0.12- $\mu$ m  $L_{eff}$  and low  $V_t$  FET's. These changes culminate into a 480-MHz RISC microprocessor at 2.0 V, 85°C, that is manufactured in a fast process.

### REFERENCES

- [1] P. Reed, M. Alexander, J. Alvarez, M. Brauer, C.-C. Chao, C. Croxton, L. Eisen, T. Le, T. Ngo, C. Nicoletta, H. Sanchez, S. Taylor, N. Vanderschaaf, and G. Gerosa, “A 250 MHz 5 W RISC microprocessor with on-chip L2 cache controller,” in *ISSCC Dig. Tech. Papers*, Feb. 1997, pp. 412–413.
- [2] D. Edelstein, J. Heidenreich, R. Goldblatt, W. Cote, C. Uzoh, N. Lustig, P. Roper, T. McDevitt, W. Motsiff, A. Simon, J. Dukovic, R. Wachnik, H. Rathore, R. Schulz, and L. Su, “Full copper wiring in a sub-0.25  $\mu$ m CMOS ULSI technology,” in *IEEE Int. Electron Device Meeting Tech. Dig.*, Dec. 1997, pp. 773–776.
- [3] J. Alvarez, H. Sanchez, G. Gerosa, and R. Countryman, “A wide-bandwidth low voltage PLL for PowerPC Microprocessors,” *IEEE J. Solid-State Circuits*, vol. 30, pp. 383–390, Apr. 1995.

**Chekib Akrouit** received the B.S. and M.S. degrees in physics and the Ph.D. degree from the University Pierre and Marie Curie, Paris, France, in 1981, 1982, and 1984, respectively.

He joined IBM, Essonnes, France, in 1984 and worked in bipolar technology and then in CMOS technology for cache macros within microprocessors. He also held several management positions developing telecommunications products. In 1994, he joined IBM, Essex Junction, VT, working on microprocessor design as Technical Lead and later as Manager. He currently is Manager of advanced PowerPC microprocessor development at IBM, Austin, TX.

**John Bialas** received the A.S. degree in electrical engineering from Lincoln Technical Institute, Allentown, PA, in 1981 and the B.S.E.E. degree from Capitol College, Laurel, MD, in 1989.

He is an Advisory Engineer specializing in very-large-scale-integration circuit design in the application-specific integrated circuit core development area. His prior activities include space system development and PowerPC advanced technology development, all with IBM. He has received five U.S. patents.



**Miles Canada** received the B.S. degree in electrical engineering from the University of Illinois, Urbana, and the B.A. degree in physics from Augustana College, Sioux Falls, SD, in 1978.

He joined IBM, Fishkill, NY, in 1979, working on emitter-coupled logic and SRAM products until 1992. Since then, he has worked on microprocessor circuit design at IBM Microelectronics, Essex Junction, VT.

**Duane Cawthron**, photograph and biography not available at the time of publication.

**James Corr** received the B.E.E. degree from the University of Minnesota, Minneapolis St. Paul, and the M.S. degree in electrical engineering from California State University.

He is an Advisory Engineer in the Advanced PowerPC Development Department, IBM, Essex Junction, VT, where he is involved in hardware design verification and test and diagnostics strategy development. Since joining IBM in 1979, he has been involved in both design and manufacturing of various microprocessors and logic devices, as well as design and development of application-specific integrated circuit design systems and system design tools. Much of this work has been in the area of automated test and diagnostics systems. His work on this new reduced instruction set computer processor has been primarily in the area of hardware verification and functional pattern diagnostics methodologies.



**Bijan Davari** (S'77–M'79–SM'92) received the M.S. and Ph.D. degrees in electrical engineering from the Rensselaer Polytechnic Institute, Troy, NY, in 1984.

He joined IBM's research division at the T. J. Watson Research Center, Yorktown Heights, NY. Since that time, he has worked on various aspects of scaled CMOS and BiCMOS technologies, including device scaling and process integration. He defined and developed a selectively scaled 0.25- $\mu\text{m}$  CMOS technology at 2.5 V, demonstrating significant

performance and power-reduction improvement over CMOS technologies at 3.3 V. This work has set the direction and the supply voltage for the post-3.3-V CMOS generations. He presently is the Director of the Advanced Logic and SRAM Technology Development Department in IBM's Semiconductor Research and Development Center, Hopewell Junction, NY. His department's activities include the development of CMOS technologies down to sub-0.1  $\mu\text{m}$  for logic and SRAM products, silicon on insulator, embedded DRAM, and nonvolatile RAM. He is the author or coauthor of more than 70 publications in various aspects of semiconductor devices and technology.

Dr. Davari is an IBM Fellow.



**Robert Floyd** received the B.S.E.E./C.S. degree from the University of Connecticut, Storrs, in 1978.

For most of his 20 years with IBM, he has worked in chip integration for microprocessors.

**Stephen Geissler** received the B.S.E.E. degree from the Rochester Institute of Technology, Rochester, NY, in 1981 and the M.Eng. degree from Cornell University, Ithaca, NY, in 1982.

He joined IBM in 1982 and is currently an Advisory Engineer responsible for phase-locked loop and high-speed clock circuit design. He has also been involved in a device design and process integration for CMOS DRAM and logic technologies. He is the author or coauthor of 12 publications and has received seven U.S. patents.

**Ronald Goldblatt**, photograph and biography not available at the time of publication.



**Robert Houle** received the B.A. degree in mathematics and the B.S. degree in electrical engineering from the University of Vermont, Burlington, and the M.S. degree in electrical engineering from the University of Connecticut at Storrs.

He joined IBM in 1980. Since 1981, he has been involved in the design of microprocessors.



**Paul Kartschoke** received the B.S.E.E. degree from Rutgers—The State University, New Brunswick, NJ, in 1990 and the M.S.E.E. from the University of Virginia, Charlottesville, in 1992.

He joined IBM in 1992 and currently is a Staff Engineer performing circuit design on a microprocessor project. He has worked on several microprocessor projects specializing in programmable logic arrays, microprocessor circuit development, and various circuit tools. He is the author of several papers and has received eight patents.

**Diane Kramer** received the B.S. degree in electrical engineering from the University of Vermont, Burlington, VT.

She joined IBM in 1988 and since has been involved in both process development and microprocessor design.



**Peter McCormick** (S'64–M'67) received the B.S.E.E. degree from Worcester Polytechnic Institute, Worcester, MA, in 1965 and the M.S.E.E. degree from Michigan State University, East Lansing, MI, in 1967.

He joined IBM, East Fishkill, NY, in 1967 and transferred to Manassas, VA, in 1970 and to Essex Junction, VT, in 1978. His work experience has been in bipolar and field-effect transistor circuit and chip design spanning SSI, LSI, and VLSI digital integrated circuits. Development projects have included master slice, master image, and custom chip designs. He currently is a Senior Engineer in the PowerPC Microprocessor Development Group of the IBM Microelectronics Division.



**Norman Rohrer** received the B.S. degree in physics and mathematics from Manchester College, North Manchester, IN, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from The Ohio State University, Columbus, in 1990 and 1992, respectively.

He has been with IBM, Essex Junction, VT, since 1992, where he works in the PowerPC microprocessor area. His interests include microprocessor development, VLSI, and technology interactions with CMOS circuit design.



**Gerard Salem** received the B.S.E.E. degree from the Polytechnic Institute of New York, Brooklyn, in 1983.

He is a Logic Design Engineer currently working on design, bringup and debug, and hardware/model critical path analysis of PowerPC processors at IBM Microelectronics, Essex Junction, VT. He joined IBM in 1983 and worked on ES/9000 processor design, recovery design, and bringup. In 1993, he joined PowerParallel and worked on RS/6000 SP communication adapter design, bringup, and test.

He has filed several patent applications.

**Lisa Su** (S'92–M'95) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge.

She joined Texas Instruments in 1994 as a Member of Technical Staff working in the SOI Process Integration group. Since 1995, she has been with the IBM Semiconductor Research and Development Center. She currently is Manager of CMOS front-end process integration.

**Linda Whitney**, photograph and biography not available at the time of publication.

**Ronald Schulz**, photograph and biography not available at the time of publication.