

A 400-MHz S/390 Microprocessor

Charles F. Webb, Carl J. Anderson, Leon Sigal, Kenneth L. Shepard, John S. Liptay, *Member, IEEE*, James D. Warnock, Brian Curran, Barry W. Krumm, Mark D. Mayo, Peter J. Camporese, Eric M. Schwarz, *Member, IEEE*, Mark S. Farrell, Phillip J. Restle, *Member, IEEE*, Robert M. Averill III, Timothy J. Slegel, William V. Huott, Yuen H. Chan, *Member, IEEE*, Bruce Wile, Thao N. Nguyen, *Member, IEEE*, Philip G. Emma, *Fellow, IEEE*, Daniel K. Beece, Ching-Te Chuang, *Fellow, IEEE*, and Cyril Price

Abstract—A microprocessor implementing IBM S/390 architecture operates in a 10 + 2 way system at frequencies up to 411 MHz (2.43 ns). The chip is fabricated in a 0.2- μm L_{eff} CMOS technology with five layers of metal and tungsten local interconnect. The chip size is 17.35 mm \times 17.30 mm with about 7.8 million transistors. The power supply is 2.5 V and measured power dissipation at 300 MHz is 37 W. The microprocessor features two instruction units (IU's), two fixed point units (FXU's), two floating point units (FPU's), a buffer control element (BCE) with a unified 64-KB L1 cache, and a register unit (RU). The microprocessor dispatches one instruction per cycle. The dual-instruction, fixed, and floating point units are used to check each other to increase reliability and not for improved performance. A phase-locked-loop (PLL) provides a processor clock that runs at 2 \times the system bus frequency. High-frequency operation was achieved through careful static circuit design and timing optimization, along with limited use of dynamic circuits for highly critical functions, and several different clocking/latching strategies for cycle time reduction. Timing-driven synthesis and placement of the control logic provided the maximum flexibility with minimum turnaround time. Extensive use of self-resetting CMOS (SRCMOS) circuits in the on-chip L1 cache provides a 2.0-ns access time and up to 500 MHz operation.

Index Terms—CMOS integrated circuits, computer architecture, integrated circuit design, logic design, microprocessors.

I. INTRODUCTION

THE single-chip microprocessor was designed for the IBM S/390 Enterprise Server Generation-4 system. This microprocessor allowed IBM to fully replace its water-cooled Enterprise ES/9000 system [1], implemented in bipolar technology, with an air-cooled system implemented in CMOS technology. The microprocessor was initially designed in IBM's CMOS5X technology and then migrated to CMOS6S technology by shrinking the FET channel length dimensions for performance but not shrinking the interconnect dimensions for time to market [2]. Typical technology parameters are shown in Table I. The technology features 0.2- μm L_{eff} , 5.5-nm gate oxide, low resistance Ti-salicided N^+ and P^+ polysilicon and diffusions, shallow trench isolation, metal fuses for laser

TABLE I
TECHNOLOGY FEATURES

L_{eff}	0.2 μm
Gate Oxide	5.5 nm
M1 Pitch	1.2 μm
M2 Pitch	1.8 μm
M3 Pitch	1.8 μm
M4 Pitch	1.8 μm
M5 Pitch	4.8 μm
Power Supply	2.5 V

blow, N^+ precision resistors, five levels of metal, and tungsten local interconnect. The power supply is 2.5 V. A die photo is shown in Fig. 1. The chip characteristics are shown in Table II. The chip is 17.35 mm \times 17.30 mm with about 7.8 million transistors. There are about 3.8 million logic transistors and 4.0 million array transistors. The measure power dissipation at 300 MHz is 37 W. There are 1600 area C4 and 448 off-chip signal I/O's. Dedicated thin-oxide capacitors [3]–[5] of 102 nF are provided for on-chip decoupling. This, combined with the "built-in" nonswitching well-to-substrate and diffusion-to-well capacitances, provides about 200-nF on-chip decoupling capacitance. The chip operated successfully in a 10 + 2 way system configuration, where the extra two processors served as I/O processors for the ten-way system. Fig. 2 shows the top view, front view, and side view of a 10 + 2 way multichip module. The module also contains eight L2 cache chips, eight bus switch node chips, and four memory bus adapter chips. The size of the module is 127.5 mm \times 127.5 mm in area and 12.4 mm thick before planarization. It consists of four thin-film wiring layers and 68 ceramic layers and handles up to 900 W (nominal) at 2.7 V with a nominal air flow rate of 150 cubic-feet per minute (CFM). It has 3526 bottom surface metallurgy (BSM) I/O's, of which 1763 are signal I/O's. In engineering system test, a 10 + 2 way system runs at

- 2.70 ns at 2.57 V with ambient air at 35°C,
- 2.62 ns at 2.70 V with ambient air at 25°C,
- 2.52 ns at 2.80 V with chiller cooling module top hat to 25°C.

A five-way system runs at 2.50 ns (400 MHz) at 2.80 V with ambient air at 10°C. A 10 + 2 way system with fast chips runs at 2.43 ns (411 MHz) at 2.50 V with chiller cooling module top hat to 12°C.

Manuscript received April 10, 1997; revised June 16, 1997.

C. F. Webb, J. S. Liptay, B. Curran, B. W. Krumm, M. D. Mayo, P. J. Camporese, E. M. Schwarz, M. S. Farrell, R. M. Averill III, T. J. Slegel, W. V. Huott, Y. H. Chan, B. Wile, and C. Price are with IBM Systems, 390 Division, Poughkeepsie, NY 12602 USA.

C. J. Anderson, L. Sigal, K. L. Shepard, J. D. Warnock, P. J. Restle, T. N. Nguyen, P. G. Emma, D. K. Beece, and C.-T. Chuang are with IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

Publisher Item Identifier S 0018-9200(97)08031-1.




Fig. 3. Pipeline for typical register-storage instruction.

The ESA/390 Floating Point Registers (FPR's) are implemented in a two-read/one-write register file. The buffer control element (BCE) contains a 64-Kbyte cache organized in 128-byte lines with a four-way set-associative absolute-address directory. The BCE also includes a 256-entry translation lookaside buffer (TLB), an eight-entry fully-associative access register translation lookaside buffer (ALB), an eight-deep store address queue, and a 64-deep 64-b-wide store data buffer. A 32-Kbyte read only store (ROS) holds frequently used internal code (millicode) routines. The cache is interleaved on a double word (8-byte) basis, and "continuation fetch" controls allow access to multiple portions of the same cache line without reaccessing the TLB and directory, significantly improving cache bandwidth for sequential instruction fetching and for long storage operands. The BCE also includes the interface to the off-chip second-level cache. The register unit (RU) maintains an error correction code (ECC)-protected copy of the architected processor state, including FR's, AR's, FPR's, ESA/390 control registers, millicode control registers, and ESA/390 timing facility. The RU also implements various system support functions, including processor error detection and recovery (to be discussed later). The pipeline for a typical register-storage instruction is shown in Fig. 3.

In order to obtain a suitable performance level, many features of the ESA/390 architecture must be implemented with hardware controls, making even a "simple" microarchitecture relatively complex. These features include a full set of register-storage and storage-storage operations, several addressing modes, byte-oriented (unaligned) storage operands, packed-decimal arithmetic, strong storage ordering and consistency (including store-in-instruction-stream), precise interruptions, and program event recording. Many hardware structures in the processor exist solely to support these features, representing a conscious tradeoff between simplicity of design and the need for high-performance execution of ESA/390 functions. Where possible, the impact to mainline execution has been confined to monitoring for unusual conditions (e.g., program exception conditions) and blocking execution of affected instructions. In these cases, the pipeline is purged and the affected instruction(s) re-executed in a nonpipelined mode or with internal code (millicode). Even this design approach placed considerable pressure on the timing of many control logic paths.

Mainframe-class reliability and availability are achieved by a unique checking and recovery design. The IU, FXU, and FPU are replicated on the chip, and all outputs that directly affect the architected processor state are sent from both copies of these units to the RU. The RU compares the copies of

the outputs and buffers the state updates for each hardware instruction. As each instruction is completed, those results are moved to a checkpoint array in which the entire architected state of the processor is maintained with ECC protection. The state updates sent to the RU from the FXU and FPU are also broadcast to the BCE and IU so that local copies of architected facilities (e.g., certain ESA/390 control registers) are maintained in lock step with changes to the master copy in the RU. To this end, the architected processor state is mapped into an 8-b address space of 32-b and 64-b registers, so that all state updates are communicated via a single address and data bus. If an error is detected in the RU comparison of the IU, FXU, or FPU outputs, or if any other error is detected in the hardware (e.g., a parity error in the cache), then updates to the checkpoint array are blocked and a CPU recovery sequence is initiated. In this sequence, completed operand stores are drained to the level-2 cache; the cache, directory, TLB, and ALB are purged; all pipeline controls in the IU, FXU, and FPU are reset; and the processor state is refreshed by reading each entry from the RU checkpoint array, passing the data through the FXU data flow, and writing it back to the same RU address, updating all IU, FXU, FPU, and BCE copies of that register in the process. This restores the processor to a consistent architected state from before the hardware fault, and instruction processing can proceed from that point. This sequence is performed entirely by hardware and is transparent both to ESA/390 code and to internal code (millicode). This design eliminates the need for error checking within the IU, FXU, and FPU logic while providing almost 100% recoverability from all transient (soft) hardware faults.

The robust checking and recovery mechanism is crucial in the mission-critical enterprise-wide server applications of S/390 systems, where customers expect and rely on "bullet-proof" design to protect their data and to support continuous operation. On the test floor, mismatches of outputs from two copies of the units have been observed during cycle time stressing and in noise-induced failures. The design also includes on-chip internal error-injection logic, which forces various errors (including mismatches of the two copies) at random and/or at controllable points in time. Well over 99% of all such random transient errors are fully recovered with no corruption of machine state. These random transient errors were either observed during cycle time stressing or intentionally injected using error-injection logic on the test floor. Fault data due to α -particle, electromigration, etc. are not available at this moment. A hardware trace facility is provided primarily for debugging logic bugs during system test but may also be used to analyze hardware errors. All errors

are logged. When an error occurs and the CPU performs its refresh/retry operation, an internal interrupt is made pending, and at the next interruptible point the millicode will get control, clean up some asynchronous interrupt controls, log out the hardware trace array, and indicate the recovery event to the service processor, which will record it in a log. No operator or software action is required. There is no option to use only one side of a failing processor since some secondary data connections are not replicated, and it is not acceptable to customers for the machine to run for even a short time with no checking. The system design does support dynamically transferring work from a failing CPU to a different CPU in the same system, including a provision for having "hot spare" CPU's on the module. There is no provision for using the redundant units for increased performance (super-scalar mode) instead of self-checking. Supporting any sort of super-scalar mode would have added far more complexity and would have required doubling most RU facilities and many BCE facilities, and would not have fit on the chip.

Some ESA/390 functions are too complex for hard wired control sequences. These functions are implemented via a form of internal code known as millicode. Millicode executes in a special mode of processor operation as a set of highly-privileged subroutines. The millicode architecture is a variant of the ESA/390 instruction set architecture and includes its own set of GR's and AR's and all of the hard-wired ESA/390 instructions. It also includes a set of special instructions (unique to millicode) which provide access to the processor state in the RU checkpoint array, directly exploit the execution hardware in the FXU, and invoke special hardware functions to accelerate performance-sensitive ESA/390 operations. Millicode is invoked when the IU decodes an ESA/390 instruction which is not implemented in hardware or when an interruption condition is detected. The FXU then executes a millicode entry operation, saving information about the ESA/390 instruction in the millicode GR's, placing the processor in millicode mode, and branching to the start of the appropriate millicode routine. The millicode routine is then executed using the same IU, FXU, FPU, BCE, and RU controls as are used for ESA/390 instruction execution. When a "millicode end" instruction is executed, the processor returns to ESA/390 mode and branches to the updated ESA/390 instruction address. Millicode instructions reside in a portion of main storage which is not accessible to ESA/390 programs. A 32-Kbyte ROS in the BCE contains frequently used millicode routines to minimize cache displacement due to millicode instruction fetches.

The emphasis on cycle time has had an impact on processor performance as measured in CPI. This is primarily due to the lengthening of the instruction pipeline (one cycle longer than in most past S/390 processors) and the use of millicode instead of horizontal microcode for complex functions. At the same time, analysis of critical timing paths in this processor and in other S/390 processor designs shows that this simpler design yields a frequency advantage that is substantially greater than the CPI disadvantage given the same technology. When compared with the previous IBM S/390 CMOS microprocessor (G3), this design achieves a cycle time improvement of 70%

(given the same technology) with 30% higher CPI, yielding a net gain of 30% in processor performance (excluding cache effects). The CPI disadvantage is somewhat larger (55%) when compared with the IBM ES/9000 model 9021 processor [1], which was a two-way S/390 super-scalar design with register renaming, out-of-sequence execution, branch target prediction, and an internal bandwidth comparable to a four-way super-scalar RISC design. Even in this comparison, the frequency advantage (approximately 60% given the same technology) is dominant, and the added benefits in chip area and design schedule further favor the simple, high-frequency architecture.

III. FLOORPLAN AND CLOCK DISTRIBUTION

The microprocessor floorplan and clock distribution is shown in Fig. 4. The BCE and RU are centrally located in the floorplan to support communication with both sets of instruction and execution units. The phase-locked loop (PLL) is located near the center of the chip and generates the internal system clocks that runs at $2\times$ the system bus frequency. It operates over the range from 36 MHz to 571 MHz with less than 4.0 ps/mV of long-term phase error between the PLL output and reference clock due to power supply noise and less than 1.0 ps/mV reduction in cycle time due to noise. A single-phase clock is distributed from the chip PLL/central clock buffer to all the latches inside the macros in three levels of hierarchy. The first two levels of clock distribution are in the form of balanced H-like trees, using primarily the top two metal layers. The first level tree routes the global clock from the central clock buffer to the nine sector buffers. Each of IU, FXU, FPU, and RU has one sector buffer while the BCE has two sector buffers. The sector buffers repower the clock to all macros inside the sectors. There are 580 macro clock pins among all the units. The clock propagation delay along the tree is balanced against macro input capacitance and RLC characteristics of the tree wires. Horizontal wiring of each tree is in low resistance Metal-5. At various places along the tree, inductive coupling is reduced and return path improved by using power wires for shielding. Decoupling capacitors are incorporated into central and sector buffers to reduce delta-I noise. A clock wiring methodology was developed with custom routing and timing CAD tools. The detailed routing as well as the widths of all clock wires were optimized to minimize skew, mean delay, power, wiring tracks, and sensitivity to process variations. Three-dimensional modeling was performed using a full-wave electromagnetic field solver [6], and distributed RLC modeling was used for virtually every wire in all the trees during the design and tuning/optimization process [7]. A number of cases were analyzed, and the results were used to generate a combination of analytic models and look-up tables containing distributed RLC parameters for all clock geometries used. Each wire segment was represented by a equivalent circuit consisting of up to six RLC " π " segments. Extensive simulations and wire width tuning [8] were done to guarantee low clock skew at macro pins. Typical simulated RLC delay of the first level tree is 300 ps with 20 ps skew at the sector buffers. The sector buffer delay is 230 ps. Typical simulated RLC delay within sectors is 210 ps with 30 ps skew

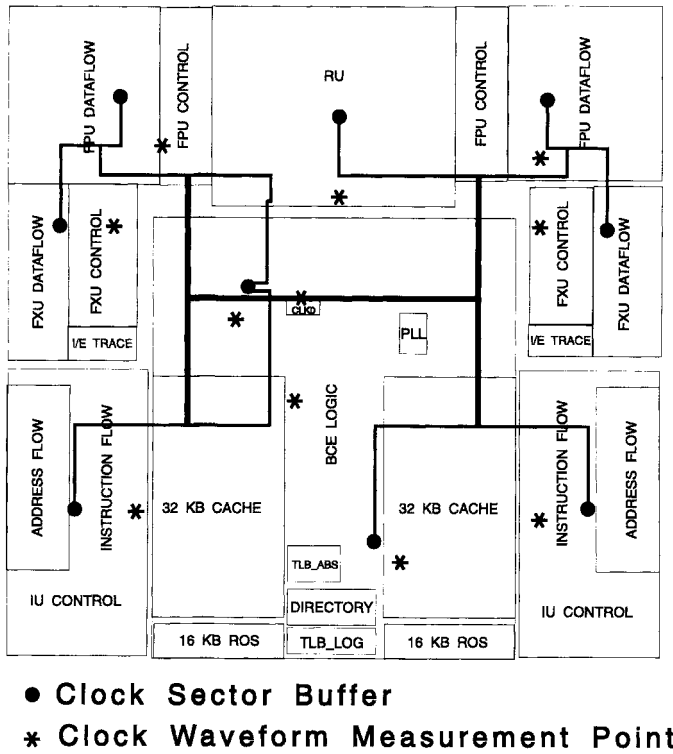


Fig. 4. Chip floorplan and clock distribution.

at the macros. The last level of clock distribution is local to each macro. Fig. 5 shows clocking scheme within macros. From the macro pin the clocks are wired to clock blocks. The overall target skew for this wire is under 20 ps. For large area macros, multiple clock pins were used to reduce wire length to clock blocks. The clock block generates local clocks that drive latches as will be explained in the next section. The target skew for local clocks is under 50 ps. All macro-level wiring is done by hand for custom macros or with a place and route tool for synthesized macros. For synthesized macros that had many latches, and therefore multiple clock blocks, a clock optimization tool was used that reassigned latches to clock blocks based on cell placement. This resulted in clock blocks driving latches that were placed closest to them. Macro layouts were extracted for *R* and *C* parasitics, and the extracted netlists were used to time the macros. This means that any skew in the last level of clock distribution was captured in that macro's timing abstraction.

Fig. 6 shows the measured waveforms of the central clock buffer output and clocks at ten points of the 580 macro pin locations (marked on Fig. 4) driven by the second level clock tree. The measurement was performed using a novel electron-beam probe with a 20-ps time resolution on the top wiring layer [9]. No special pads or test structures were required to enable the electron-beam probing, since the clock distribution used top metal layer for most horizontal wiring, and the spatial resolution of the probe (2.0 μm) allowed probing of minimum width wires. The only chip preparation needed was the removal of area C4 solder balls and all passivation above the top wiring level. Because the chip was powered using a standard cantilever probe card in the electron-beam probe, the chip

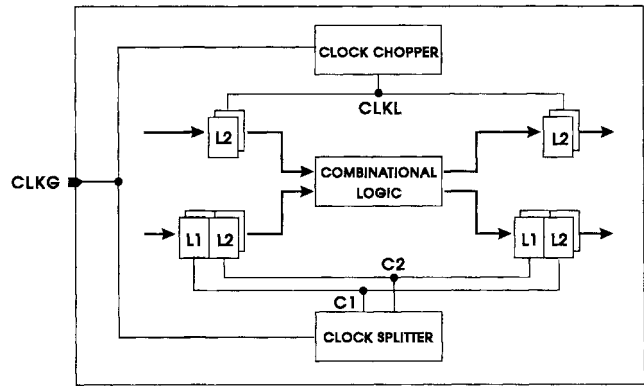


Fig. 5. Clocking scheme within macro.




Fig. 6. E-beam measured clock waveforms at macro pin locations marked on Fig. 4.

clock was run at low frequency to reduce power supply noise. Power supply noise during these measurements was measured to be less than 100 mV. The results indicate a mean delay of 740 ps and less than 30 ps skew from the central clock buffer to the macro pins.

IV. CLOCK BLOCKS AND LATCHES

There are two types of latches used in the microprocessor outside the arrays: L2 only latch and L1-L2 pair. Both latch types are cycle boundary latches (i.e., there are no midcycle or phase latches). Cycle boundary is defined to occur at "CLKG" falling. Corresponding to the two latch types are two types of clock blocks. The first clock block/latch combination is shown in Fig. 7(a). This clock block chops the global clock on the falling edge to create a short pulse "CLKL" that triggers the latch [10]. By using either a dynamic multiplexer [Fig. 7(a)] or a preset static multiplexer [Fig. 7(b)] in front of the latch, a fast latch delay is achieved. This mux/latch combination interfaces smoothly with the static circuits and yet allows fast delays for multi-input high-fanout registers typical of the data flow. While "CLKL" is inactive (high), the dynamic multiplexer is in the precharge state while the latch is holding its state. Similarly, the static multiplexer is preset high (node mux_A is high) while the latch is holding its state. When "CLKL" low going pulse arrives, the latch node "LAT" begins to discharge. If multiplexers' data/selects are such that the multiplexers remain in their precharged/preset state, node "LAT" fully discharges. If, however, the multiplexers evaluate, then node "LAT" is driven high. The n/p transistors in the multiplexers

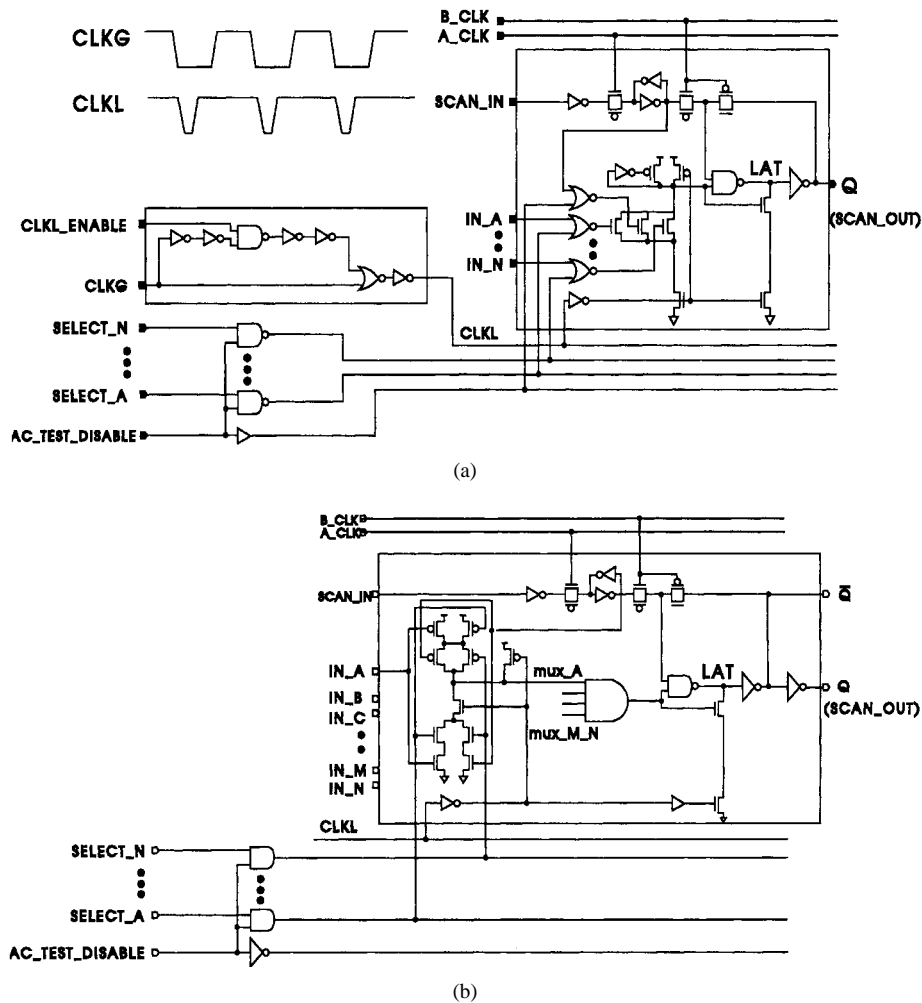


Fig. 7. Clock block/latch combination: (a) generation of local clock "CLKL" and latch with dynamic multiplexer and (b) latch with preset static multiplexer.

are skewed to favor the transition launched by the arrival of the "CLKL" pulse. Strength ratios in the preset static multiplexer are limited to 2.5:1 to maintain reasonable noise margin and allow adequate time for presetting all gates at the end of the clock pulse. The seven-input preset static multiplexer evaluates in about 225 ps compared to about 400 ps for a standard static multiplexer with the same input capacitance and area. When "CLKL" is active the latch is in the transparent mode.

The second clock block/latch combination is shown in Fig. 8. This clock block splits the global clock on the falling edge to create C1/C2 clocks. C1-C2 clock-overlap at the cycle boundary is set close to 0 ps. Having a positive overlap would reduce the latch propagation delay but it would also require more early-mode padding. These latches are used in nontiming-critical data flow macros and in control macros where all latches are single input and the speed advantage of an L2-only latch is reduced.

All latches used in the design are fully scannable and level sensitive scan design (LSSD) compatible. The overall chip area penalty is less than 5%. Having an LSSD-compatible design increased productivity during test vector generation and allowed > 99.5% dc (stuck at fault) test coverage. AC test coverage often suffers from the inability to create appropriate transitions due to latch adjacency in the scan chain. This is

most prevalent in dataflow macros with complicated logic functions (e.g., multipliers). To eliminate the latch adjacency problems, nonfunctional scan-only latches were inserted into some macros for every register bit. With this addition, > 91% ac transition test coverage was achieved.

Special circuitry was added to all clock blocks to allow edge-shifting at the cycle boundary. Taking the clock splitter as an example, we are able to do the following.

- Completely unoverlap C1 falling and C2 rising edges by large amount. This provides a work-around for potential early-mode problems.
- Delay C1 falling from its nominal value. This allows stressing of early-mode and provides a determination of how much margin exists in the design.
- Delay C1 falling and C2 rising together from their nominal values. This allows cycle-stealing from the previous cycle.

The last feature was found to be extremely useful in debugging late-mode problems. To facilitate clock stressing and late-mode debugging, each macro contains a 3-b scan-only register, which controls all the clock stressing functions in the clock blocks located in the macro. Since the clock stressing is local to each macro, when the first late-mode path was found, more

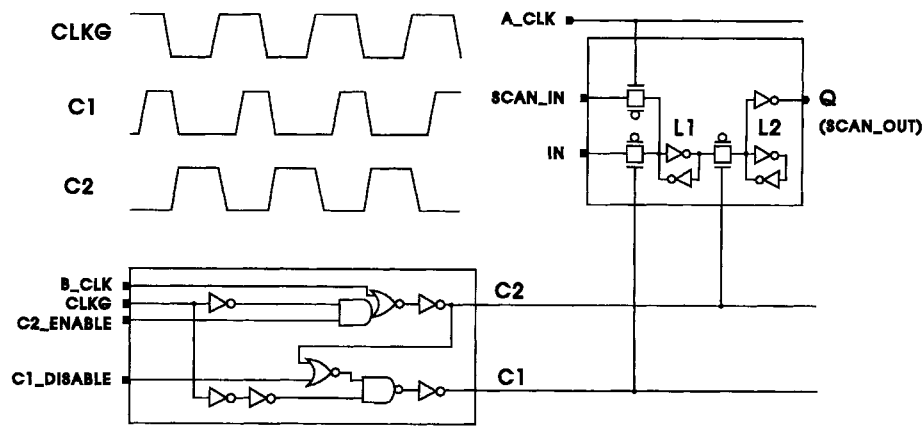


Fig. 8. Clock block/latch combination: generation of C1/C2 clocks for L1/L2 latches.

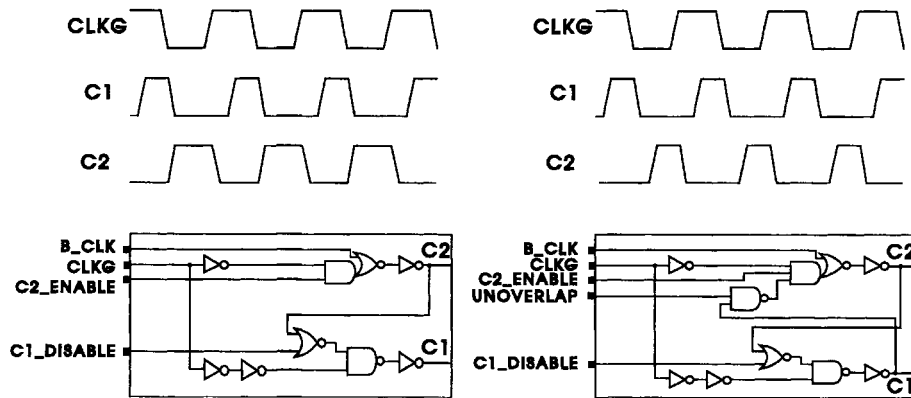


Fig. 9. Use of clock-splitter/clock block to separate C1 falling edge and C2 rising edge.

time was given to the failing cycle through cycle-stealing. This usually did not create late-mode problems in the following cycle for the stressed macro. With the worst late-mode path now masked, work could continue in finding the next late-mode path that failed. Fig. 9 shows how clock-splitter/clock block was modified to cause unoverlap between C1 falling and C2 rising edges.

V. GLOBAL DESIGN STYLE AND ISSUES

The most fundamental strategy was to design the dataflow to the most aggressive cycle time achievable and then match the control portion to the dataflow cycle time. In cases where the control paths lag the dataflow cycle time after all possible logic optimization, the CPI was sacrificed for cycle time. A concurrent bottom-up and top-down design approach was adopted for achieving high-frequency operation as well as for time-to-market. The dataflow are full-custom bottom-up designs with mostly static circuits except the dynamic multiplexer. The control portions are synthesized with top-down macro aspect-ratio and pin placement and placed/routed using static books with fine power level granularity. Two carefully tuned standard cell booksets were developed for use with synthesis. One of these booksets contained fixed schematics and layouts and the other bookset was parameterizable with capabilities for automatic layout generation. Both bookset's designs were limited to relatively simple functions. The unit and macro interconnect were implemented in parallel to the macro layout.

The power distribution supports an average dc voltage drop of 23 mV. The Delta-I current transients were managed by including additional on-chip decoupling capacitors around large noise sources such as the off-chip drivers, clock buffers, and on-chip drivers with large loads. Since a large amount of switching capacitance occurs in the dataflow stacks, decoupling capacitors were also placed under the wiring tracks. Dedicated thin-oxide capacitors of 102 nF are provided for on-chip decoupling [3]–[5]. This, combined with the “built-in” nonswitching well-to-substrate and diffusion-to-well capacitances, provides about 200 nF on-chip decoupling capacitance. The thin-oxide capacitor features a “built-in” fuse mechanism where “weak” spots between M1 and contact are used to blow connections to V_{DD} and GND in the presence of large current resulting from oxide defects. Each capacitor also has a gated NFET control device with an external “decap_enable” pin for leakage current measurement during test. A special decoupling capacitor cell (Fig. 10) is designed to fit under the dataflow wiring tracks. The cell is double bit-pitch wide ($43.2\ \mu\text{m}$) and 14 tracks tall ($25.2\ \mu\text{m}$). Two out of the 14 horizontal wiring tracks are specifically blocked for the decoupling capacitor wiring so the capacitor can fit right under the wiring tracks. A low-resistance layout of the capacitor cell provides a fast time constant of about 85 ps.

The 64-kbyte unified cache features a $33.2\text{-}\mu\text{m}^2$ planar six-transistor cell. Fast signal conversion between static CMOS and self-resetting CMOS (SRCMOS) circuits, and extensive



Leon Sigal received the B.S. degree in biomedical engineering in 1985 from the University of Iowa, Iowa City, and the M.S. degree in electrical engineering in 1986 from the University of Wisconsin, Madison.

He worked at Hewlett Packard's Microprocessor Development Laboratory between 1986 and 1992. He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1992, and is currently a Research Staff Member. He has been leading the circuit design for the CMOS S/390 microprocessor.



Brian Curran received the B.S. degree in electrical engineering from the University of Wisconsin at Madison in 1984.

He joined the IBM Data Systems Division (now System/390 Division) that same year. He has designed logic and circuits for large system processors and memory subsystems. He was a member of the teams which developed the System/3090 and ES/9121 processors. Most recently, he lead the high-frequency design effort of a 400-MHz CMOS S/390 microprocessor. He has been issued seven patents relating to processor design and has three patents pending. He is currently a Senior Engineer in the IBM System 390 Division and a technical leader in the development of a future high-frequency full-custom CMOS microprocessor.

Mr. Curran is the recipient of several IBM Technical Achievement Awards.

Kenneth L. Shepard received the B.S.E. degree in electrical engineering from Princeton University, Princeton, NJ, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1992 and is presently a Research Staff Member in the VLSI Design Department. He is the technical leader for the design methodology for the CMOS S/390 microprocessor. His present interests are in computer-aided design of VLSI systems. He is also an Adjunct Assistant Professor of Electrical Engineering at Columbia University, New York, NY.

Dr. Shepard received the Hertz Foundation doctoral thesis prize for his research on nanometer-scale device physics.



Barry W. Krumm received the B.A. degree in computer science modified with electrical engineering from Dartmouth College, Hanover, NH, in 1980.

He joined IBM Data System Division, Poughkeepsie, NY, in 1980, and has since worked on the logic design of buffer control element and storage control element. He is currently an Advisory Engineer in the IBM System 390 Division, leading the Buffer Control Element design team for CMOS S/390 microprocessor.



John S. Liptay (S'63-M'64) received the B.E.E. degree in 1962 and the M.E.E. degree in 1966, both from Rensselaer Polytechnic Institute, Troy, NY.

He joined IBM in 1965 at the T. J. Watson Research Center, Yorktown Heights, NY, transferring shortly thereafter to the IBM Product Development Laboratory in Poughkeepsie, NY, where he has remained since and is now a Senior Technical Staff Member. He has worked on the System/360 Models 65 and 85, the System/370 Model 168, the IBM 3033, the high-end ES/9000 processor, and the

System/390 G4 microprocessor, all in the area of CPU design. He holds 11 patents.

Mr. Liptay is a member of Tau Beta Pi, Eta Kappa Nu, and the ACM.



Mark D. Mayo received the B.S.E.E. degree with high honors from Lehigh University, Bethlehem, PA, in 1980 and the M.S.E.E. degree from Syracuse University, Syracuse, NY, in 1985.

He joined IBM, East Fishkill, NY, in 1980, and has since worked in the design of bipolar and CMOS gate arrays, standard cells, and custom microprocessor circuits. He is currently an Advisory Engineer in the IBM System 390 Division and a technical leader in the circuit design of the buffer control element (BCE) for CMOS S/390 microprocessor.



James D. Warnock received the B.S. degree in physics from the University of Ottawa, Ottawa, ON, Canada, in 1980 and the Ph.D. degree in physics from the Massachusetts Institute of Technology, Cambridge, in 1985.

He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1985, spending two years on solid-state physics research using femto-second optical techniques. Since then, he has spent time in research on silicon technology focusing on advanced bipolar, CMOS, and BiCMOS processes. He joined

the VLSI Design Department in 1993 and is presently involved in CMOS circuit design for high-performance microprocessors. He is a Research Staff Member and a Technical Team Leader for the instruction unit circuit design of the CMOS S/390 microprocessor.



Peter J. Camporese received the B.S. degree in electrical engineering from Polytechnic University in 1988 and the M.S. degree in computer engineering from Syracuse University, Syracuse, NY, in 1991.

He joined the IBM Data Systems Division, Poughkeepsie, NY, in 1988, where he has worked on system performance, circuit design, physical design, and chip integration. He is currently an Advisory Engineer in the IBM System 390 Division and a technical leader for the physical design and

integration of CMOS S/390 Microprocessor.



Eric M. Schwarz (M'84) received the B.S. Eng. Sc. degree from the Pennsylvania State University, State College, in 1983, the M.S.E.E. degree from Ohio University, Athens, in 1984, and the Ph.D.E.E. degree from Stanford University, Stanford, CA, in 1993.

Since 1984, he has been employed by IBM Corporation, first in Endicott, and now in Poughkeepsie, NY, in their mainframe division. His responsibilities include designing the arithmetic algorithm, proving the correctness, and implementing the logic design.

He is a technical leader for the fixed and floating point units for the S/390 CMOS microprocessors. He is author of nine issued patents and several articles in journals and conference proceedings. His current research interests are in computer arithmetic and architecture.



Timothy J. Slegel received the B.S.E.E. and M.S.E.E. degrees from Lehigh University, Bethlehem, PA, in 1980 and 1982, respectively.

He joined IBM in 1982 and worked on the logic design of IBM's vector processors. He has held technical positions in a number of areas including floating-point unit design, systems architecture, cache design, and CAD tool development/programming. He is currently a Senior Engineer in the IBM System 390 Division, Poughkeepsie, NY, working on the design of a

future large-system CMOS microprocessor.



Mark S. Farrell received the B.S. degree in electrical engineering from the University of Pittsburgh, PA, in 1977.

He has worked at IBM, Poughkeepsie, NY, since 1977 in the processor microcode area on many different projects. He has many patents on microcode/hardware and architecture design.



William V. Huott received the B.S. degree in electrical engineering from Syracuse University, Syracuse, NY, in 1984.

He joined IBM in 1983 at the East Fishkill, NY, facility and worked on the design/testing/debugging of high-speed logic gate arrays and memory arrays. He is currently an Advisory Engineer in IBM System 390 Division, Poughkeepsie, NY, working on custom CMOS microprocessor design with responsibilities in the definition and design of the array-built-in-self-test (ABIST) circuitry and interfaces

for high-frequency custom microprocessors. He is also responsible for the engineering test and debug of high-frequency custom micros and support chips. He holds four U.S. patents and has contributed to several publications.



Phillip J. Restle (M'86) received the B.A. degree in physics with honors in 1979 and the M.S. and Ph.D. degrees in physics from the University of Illinois, Urbana, IL, in 1982 and 1986.

In 1986, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, and worked on CMOS device measurements and modeling, oxide trap noise in small area FET's, low-temperature CMOS, multichip package testing, and DRAM testing. For the past four years, he has been in the VLSI Design Department, creating and using tools

to design clock distributions for IBM's high-performance microprocessor and ASIC chips.

Dr. Restle received an IBM Fellowship in Physics in 1984.



Yuen H. Chan (S'75-M'77) received the B.S. degree in electrical engineering from Union College in 1977 and the M.S.E.E. degree from Syracuse University, Syracuse, NY, in 1984.

He joined IBM at the East Fishkill, NY, facility in 1977 and has worked on high-performance bipolar, biCMOS, and CMOS logic and array development. He is currently a Senior Engineer at IBM System 390 Division, Poughkeepsie, NY, working on custom VLSI circuits and SRAM designs. He is a Technical Team Leader of the CMOS S/390

microprocessor array design team. He has authored and co-authored a number of technical papers and holds 15 U.S. patents.

Mr. Chan has received an IBM Outstanding Technical Achievement award and seven IBM Invention Achievement Awards.



Robert M. Averill III received the B.S.E.E. degree from Northwestern University, Chicago, IL, in 1983 and the M.S.E.E. degree from Syracuse University, Syracuse, NY, in 1990.

He joined IBM, East Fishkill, NY, in 1983 in the Advanced Test Systems Development Group. He was the Technical Lead Engineer for the Pin Electronics Group, responsible for designing 600-MHz pin electronics for automated test equipment. Since 1993, he has been working at IBM System 390 Division, responsible for the physical design

of the floating point multiplier and floating point controls of CMOS S/390 microprocessor. He is currently the integration leader for the next generation IBM CMOS S/390 microprocessor.



Bruce Wile received the B.S. degree in computer science from the Pennsylvania State University, State College, in 1984.

He joined IBM in 1985 and has worked in verification since. His verification experiences included storage controller element simulation for the S/390 bipolar ES/9000 machines including the six-way, eight-way, and ten-way multiprocessor systems. He was previously Verification Team Leader for the ten-way ES/9000 system. He is currently a Senior Engineer in IBM System 390 Division, Poughkeepsie, NY, and the Verification Team Leader for the S/390 G4 microprocessor.

Mr. Wile received an IBM Excellence Award in 1992, a Team Award in 1993, and an IBM Invention Achievement Award in 1995 for inventions and patent submissions pertaining to the Timediag and Genrand tool set.

Thao N. Nguyen (S'75-M'82) received the B.S.E.E. degree with first class honors and university medal from the University of New South Wales, Australia, in 1978, the M.S.E.E. and Ph.D.E.E. degrees from Stanford University, Stanford, CA, in 1980 and 1984, respectively. His graduate work involved submicrometer MOSFET device physics, modeling, fabrication, and characterization.

Since joining the IBM T. J. Watson Research Center in 1983, he has worked on a wide range of R&D projects and held various technical and management positions. During the first six years, he has performed and managed research in low thermal-budget semiconductor processing technologies including UHV/CVD epitaxy, PECVD insulators, integrated vacuum processing, rapid-thermal processing, ultra-thin gate dielectrics as well as electrical characterization techniques. He spent the next two years in product development at IBM East Fishkill Semiconductor Facilities, first as an Executive Technical Assistant and later as Manager of Process Integration. From 1991 to 1996 he participated in the development of an advanced RISC microprocessor for RS/6000 and a high-performance microprocessor for S/390 systems. He was a Team Leader responsible for floorplanning and chip integration for the S/390 MPU project. His current research activities are focused on the software for electronic commerce on the Internet, particularly the performance and scalability of web servers, databases, and transaction servers. He has authored or co-authored more than 40 papers and has been awarded several patents in silicon and circuit technologies.

Dr. Nguyen is a recipient of several IBM Research Division Awards, a Technical Group Award, and an Outstanding Technical Achievement Award. He has served on the Technical Program Committee of several IEEE conferences including Symposium on VLSI Technology and Symposium on VLSI Circuits.



Philip G. Emma (M'83-SM'92-F'97) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana Champaign.

In 1983, he joined the Advanced Computer Architecture group at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he did research in high-performance computer architecture. In 1992, he became the Manager of VLSI Systems and Design and was a Technical Team Leader on a CMOS S/390 microprocessor. He is currently the

Technical Assistant to the Vice President of Systems Technology and Science at the IBM T. J. Watson Research Center.

Dr. Emma has received 18 IBM Invention Achievement Awards and has earned the IBM Research Division's honorary title of Master Inventor.



Daniel K. Beece received the B.S. degree in engineering physics from Cornell University, Ithaca, NY, and the M.S. and Ph.D. degrees in physics from the University of Illinois, Urbana Champaign.

He is currently a Research Staff Member in the Computer-Aided Design and Verification Group at the IBM T. J. Watson Research Center, Yorktown Heights, NY. His research interests include design methodologies, simulation, synthesis, and verification.



Ching-Te Chuang (S'78-M'82-SM'91-F'94) received the B.S.E.E. degree from National Taiwan University, Taipei, Taiwan, in 1975 and the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1982.

He joined the Bipolar Devices and Circuits Group at IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1982, working on scaled bipolar devices, technology, and circuits. He studied the scaling property of epitaxial Schottky barrier diodes, did pioneering works on the perimeter effects of advanced double-poly self-aligned bipolar transistors, and designed the first sub-nanosecond 5-Kb bipolar ECL SRAM. From 1986 to 1988, he was manager of the Bipolar VLSI Design Group, working on low-power bipolar circuits, high-speed high-density bipolar SRAM's, multi-Gb/s fiber-optic data-link circuits, and scaling issues for bipolar/BiCMOS devices and circuits. Since 1988, he has been Manager of the High Performance Circuit Group, investigating high-performance logic and memory circuits. Since 1993, his group has worked primarily on the circuit design of high-performance CMOS microprocessors. He holds six U.S. patents and has authored or co-authored more than 110 papers.

Dr. Chuang served on the Device Technology Program Committee for IEDM in 1986 and 1987, and the Program Committee for Symposium on VLSI Circuits from 1992 to 1997. He was Publication/Publicity Chairman for Symposium on VLSI Technology and Symposium on VLSI Circuits in 1993 and 1994. He was elected an IEEE Fellow in 1994 "for contributions to high-performance bipolar devices, circuits, and technology." He is a member of the New York Academy of Science.



Cyril Price graduated from Pratt Institute, Brooklyn, NY, with the B.E.E. degree in 1969. He received the M.S.E.E. and D.E.E. degrees from Syracuse University, Syracuse, NY, in 1977 and 1986, respectively.

He joined IBM in 1969 as a Circuit/Chip Designer and then served two years as a U.S. Army Signal Corps officer. In 1978, he was named Manager of Circuit and Subsystem Design. From 1979 to 1981, he was Manager of LSI Design and Manager of Exploratory Circuits. From 1981 to 1992, he held several management positions including Advanced Design Manager responsible for CMOS chip development and IBM Kingston Technology Manager responsible for technology support of IBM Enterprise System/9000 Type 9121 air-cooled processors. In 1992, he was named Manager of Advanced Development responsible for the initial research and development and successful product introduction of IBM's S/390 G4 microprocessor. He is currently a Senior Engineer Manager in the Product Design organization working on the development of IBM's S/390 CMOS servers in Poughkeepsie, NY.

Dr. Price is a member of Tau Beta Pi, Pi Mu Epsilon, Eta Kappa Nu, and the Society of American Military Engineers (SAME). Among his achievements are an IBM Outstanding Technical Achievement Award for "Timing Driven Placement" and an Invention Achievement Award for patents and technical disclosures.