

**FA 10.3: A 400MHz S/390 Microprocessor**

Charles F. Webb, Carl J. Anderson<sup>1</sup>, Leon Sigal<sup>1</sup>, Kenneth L. Shepard<sup>1</sup>, John S. Liptay<sup>1</sup>, James D. Warnock<sup>1</sup>, Brian Curran, Barry W. Krumm, Mark D. Mayo, Peter J. Camporese, Eric M. Schwarz, Mark S. Farrell, Phillip J. Restle<sup>1</sup>, Robert M. Averill III, Timothy J. Siegel, William V. Huott, Yuen H. Chan, Bruce Wile, Philip G. Emma<sup>1</sup>, Daniel K. Beece<sup>1</sup>, Ching-Te Chuang<sup>1</sup>, Cyril Price

IBM Systems 390 Division, Poughkeepsie, NY

<sup>1</sup>IBM T. J. Watson Research Center, Yorktown Heights, NY

A microprocessor implementing IBM S/390 architecture operates in a system at up to 400MHz (2.5ns). The microprocessor, initially in IBM CMOS5X technology, migrated to CMOS6S by shrinking the FET length dimensions but not shrinking the interconnect dimensions (Table 1). The chip is 17.35x17.30mm<sup>2</sup> with about 7.8M transistors (Figure 1). The power supply is 2.5V and measured power dissipation at 300MHz is 37W. The microprocessor features two instruction units (IU), two fixed point units (FXU), two floating point units (FPU), a buffer control element (BCE), and a register unit (RU). It dispatches one instruction per cycle. A PLL provides a processor clock at 2x the system bus frequency.

The IU handles instruction fetch, instruction decode, address generation, and operand fetch functions. The FXU is a 64b dataflow stack and maintains the condition code and controls the taking of interrupts. A single register file (5-read/1-write) is used by the IU for address generation and by the FXU for execution. The FPU contains a radix-8 Booth-encoded multiplier. Most floating-point instructions are pipelined 1-per-cycle with a latency of 3 cycles. The FPU also executes division and square root, and extended precision and fixed point multiply and divide instructions. The BCE contains a unified 64kB L1 cache organized in 128B lines with a 4-way set-associative absolute-address directory. The cache is interleaved on a double word basis. The RU maintains an ECC-protected copy of the architected processor states. The RU also implements system support functions, including processor error detection and recovery. The pipeline for a typical register-storage instruction is shown in Figure 2.

Some ESA/390 functions too complex for hardwired control sequence are implemented via a form of internal code known as millicode. Millicode instructions reside in a portion of main storage not accessible to ESA/390 programs. A 32kB ROM in the BCE contains frequently-used millicode routines to minimize cache displacement due to millicode instruction fetches.

The IU, FXU, and FPU are replicated on 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. As each instruction is completed, the results are moved to a checkpoint array in which the entire architected state of the processor is maintained with ECC protection. If an error is detected in RU comparison, updates to the checkpoint array are blocked and a CPU recovery sequence is initiated. This design eliminates error checking within IU, FXU, and FPU while providing almost 100% recoverability from all transient (soft) hardware faults.

The floorplan and clock distribution are shown in Figure 3. A single clock is globally distributed from the chip PLL/central clock buffer to all the macros in two-levels of balanced H-like trees. The first-level tree routes the global clock from the central clock buffer to 9 sector buffers. The sector buffers repower the clock to 580 macro pins within the units. Typical calculated RLC delay of the

first-level tree is 300ps with 20ps skew at the sector buffers. The sector buffer delay is 230ps. Typical calculated RLC delay in sectors is 210ps with 30ps skew at the macros. Figure 4 shows measured central clock output and clocks at 6 points of the 580 macro pin locations (marked on Figure 3) driven by the second-level clock tree. The results indicate a mean delay of 740ps and <50ps skew from the central clock buffer to the macro pins.

At the macro, the global clock goes through the clock block. There are two types of clock blocks. The first clock block/latch combination is shown in Figure 5. This clock block chops the global clock on the falling edge to create a short 'CLKL' that triggers the latch. By using either a dynamic multiplexer or a preset static multiplexer in front of the latch, the mux/latch combination interface smoothly with the static circuits and yet allows fast delays for multi-input high-fanout registers typical of the dataflow (Figures 5, 6). For the preset static multiplexer, the output of the first mux stage is preset to high when the clock is inactive, thereby presetting all the subsequent logic stages up to the latch. The n/p transistor widths are skewed to favor the transition launched by the arrival of the clock pulse. A 7-way preset static mux evaluates in 225ps, compared to 400ps for a standard static mux with the same input capacitance and area.

Full-chip macro/wire RC extraction is for both late mode and early mode. Extra circuitry added to the clock chopper (1) delays leading edge of the pulse to allow cycle stealing, and (2) delays trailing edge of the pulse for early mode stressing.

The second clock block/latch combination is shown in Figure 7. This clock block splits the global clock on the falling edge to create C1/C2 clocks. The L1-L2 latches require less signal padding to protect against early mode. These latches are used in non-timing-critical dataflow macros and in control macros where the latches are single-input and the speed advantage of type-one latch is reduced. Extra circuitry added to the clock splitter delays C1 falling and C2 rising edges for cycle stealing and early mode stressing. Both latch types are LSSD compatible, and >99.5% dc and >91% ac test coverage are achieved.

Dedicated 102nF thin-oxide capacitors provide on-chip decoupling. A special decoupling capacitor cell with built-in fuse mechanism and 120ps time constant fits under the dataflow wiring tracks. Dataflow designs are full-custom with mostly static circuits except the dynamic multiplexer. Control portions are placed and routed using static books with fine power granularity. The 64kB cache features a 33.2µm<sup>2</sup> planar 6-T cell with ABIST capability. Extensive use of SRCMOS circuits achieve 2.0ns access and up to 500MHz operation [1].

*Acknowledgment:* See page 449

#### Reference:

[1] Pelella, A., et al., "A 2ns Access, 500MHz 288Kb SRAM Macro," Dig. Tech. Papers, Symp. VLSI Circuits, pp. 128-129, 1996.

Leff	0.2µm
Gate oxide	5.5nm
M1 pitch	1.2µm
M2 pitch	1.8µm
M3 pitch	1.8µm
M4 pitch	1.8µm
M4 pitch	4.8µm
Power supply	2.5V

**Table 1: Technology features.**

Figure 1: See page 449.




Figure 2: Pipeline for typical register-storage instruction.




Figure 4: Measured clock waveforms at macro pin locations marked on Figure 3.

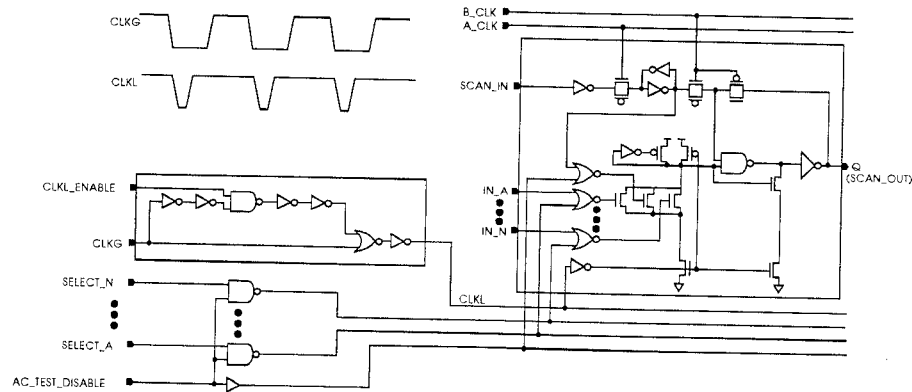


Figure 5: Clock block/latch combination: generation of local clock 'CLKL' and latch with dynamic multiplexer.

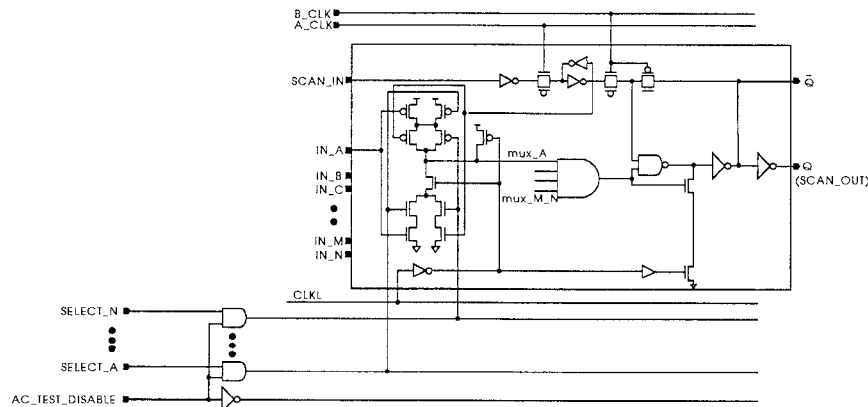


Figure 6: Clock block/latch combination: latch with preset static multiplexer.

Figure 7: See page 449.

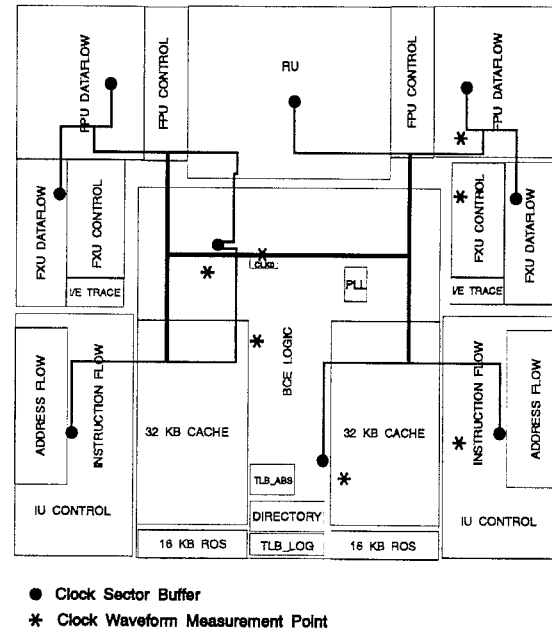




Figure 3: Chip floorplan and clock distribution.



(Continued from page 169)



*Acknowledgment:*

The authors acknowledge contributions of R. Allen, J. Badar, D. Bair, U. Bakhru, D. Balazich, K. Barkley, L. Bendrihem, M. Billeci, F. Bozeo, J. Braum, T. Bucelot, B. Bunce, J. Burns, C. Bui, S. Carey, Y. Chan, M. Check, C. Chen, K. Chin, E. Cho, G. Ditlow, K. Eng, R. Franch, J. Feldman, B. Giameli, J. Gilligan, B. Grossman, R. Hanson, Jr., A. Hartstein, R. Hatch, D. Heidel, D. Hillerud, D. Hoffman, K. Jenkins, J. Ji, G. Jordy, D. Knebel, A. Kohli, N. Kollesar, T. Koprowski, S. Kowalczyk, C. Krygowski, L. Lacey, L. Lange, K. Lewis, W. Li, P. Liu, P. Lu, J. Ludwig, V. Lund, S. McCabe, T. G. McNamara, T. McPherson, D. Merrill, P. Minear, M. Mullen, J. Navarro, J. Neely, H. Ngo, T. Nguyen, G. Northrop, D. Ostapko, P. Patel, A. Pelella, E. Pell, III, J. Rawlins, W. Reohr, S. Risch, B. Robbins, R. Robortaccio, G. Scharff, M. Scheuermann, A. Sharif, K. Shum, H. Smith, P. Strenski, S. Swaney, F. Tansi, G. Tepapas, A. Tuminaro, G. VanHuben, S. Walker, L. Wang, D. Webber, P. Williams, B. Winters, T. Wohlfahrt, D. Wong, B. Wu, P. Wu, F. Yee.