FP 13.1: A Quad-Issue Out-of-Order RISC CPU

Jon Lotz, Gregg Lesartre, Samuel Naffziger, Don Kipp

Hewlett Packard Company, Ft. Collins, CO

A 64b 4-way superscalar PA-RISC microprocessor systems operating from at 150-250MHz combines full out-of-order execution with low-cycle time, to produce >360specint and >550specfp. Specialized latching and clock circuits and extensive use of dynamic logic enable high frequencies. 3.8M logic transistors are integrated on a 17.68x19.1mm^2 die in 3.3V 0.5µm CMOS. Additional specifications are summarized in Tables 1 and 2. Figure 1 is a chip block diagram, and Figure 2 is a chip micrograph.

Improved architectural efficiency of the CPU is accomplished by combining dynamic instruction scheduling with a highly parallel 4-way superscalar execution unit [1]. On the memory side, a 64b, 768MB/s split-transaction system bus supplies the CPU with instructions and data and provides for multiprocessing [2]. The result is approximately two specint92/MHz efficiency.

The second goal for the chip is high-frequency operation. To achieve this, aggressive circuit design is used with a high performance clocking system [3]. The result is a chip/cache system capable of running up to 250MHz under nominal operating conditions (room temperature, 3.3V).

One primary feature of this chip isa large instruction reorder buffer (IRB). This 56-entry buffer is capable of reordering instructions on the fly, and tracking all the dependencies between instructions [4]. In addition to tracking dependencies, the IRB also tracks branch prediction status and is capable of flash invalidating instructions from the IRB that are incorrectly fetched. Fetching down the correct path is then resumed without any wasted patch-up cycles. The IRB comprises nearly 0.85M of the transistor count and approximately 20% of the die area.

A second major block on the chip is the floating point unit. This block is capable of starting 2 multiply accumulates (MACs) per cycle and has a three-cycle latency. To achieve this latency at 250MHz, extensive use of dynamic logic is required [5].

Fundamental to high-speed design is a low-skew, high-performance clock distribution network. This problem is solved by using a 3-level H-tree-based clocking network. A central first-level buffer generates differential clock signals that are routed in a balanced H-tree to 12 second-level buffers distributed around the chip. The resulting clock network (CK) is strapped together with wide metal-5. This keeps distribution and RC skews under 170ps across the die.

The third-level buffers or gaters incorporate high-gain and simple logic to qualify the local clocks (Figure 3). By allowing logic to be performed in the gater, a qualified, synchronous clock can be used to control register sets and dumps. This simplifies register design and reduces area. Close to 7k gaters are used in the CPU.

For any high-frequency superscalar machine providing single-state operation latencies, operand bypassing is a challenge. The operands must be provided quickly into the appropriate functional unit whether they come from the general registers, or are bypassed from any of a large number of functional units. The out-of-order nature of this machine adds the requirement of also bypassing quickly from the large rename register file. Direct bypassing from functional unit to functional unit is handled without dedicated on-the-fly comparators. Instead the IRB that already must track the launching of dependent operations, remembers when the dependency cleared and signals the bypass logic in the data stack when a bypass is required from a given functional unit.

When an instruction is inserted into the IRB, it records which rename register it should get a given operand from, or whether it should instead get its data from the general registers. By the time the instruction launches, however, the result of the producing instruction may have been moved from the rename registers to the general registers. The IRB cannot track when this happens, so the read logic for the rename and general register files must figure out, on the fly, which should provide the requested operand. This is accomplished by tagging the rename registers with one additional color tag bit to indicate whether the requested data came from an even or odd iteration through the IRB. The color bit for a given rename register is toggled whenever an instruction is retired from that register. On launch of an instruction the full tag is sent to the register decode logic in dual rail form. The extra color bit, from the time the instruction is inserted, is compared to the current color bit in the targeted rename register. A dynamic match line is generated and forwarded to the general register decode block. If the color match, and the tag is indicated to be valid, the rename registers are allowed to dump the appropriate register to the operand bus. If a match of the color bit is not indicated or the tag is invalid, the general registers are allowed to dump. (Figure 4)

Special attention is also paid to register dump circuits since their half-cycle access budget is directly in the execution path and they are replicated 64x88x2 times. A 9 read port, 4 write structure minimizes bit line discharge time and wordline capacitance by using a NOR dump structure tied to a precharge pulldown bus. In addition, a pass-gate structure for word line generation minimizes load on the dump clock (Figure 5). For the high fanout 88 register dump, a bi-directional clocked buffer quarters RC delay and reduces capacitance seen by any one dumper. The delay through this buffer is more than offset by the reduction in bit line RC. Total area for the 88x64 bit register file is 2560x2220µm^2.

Acknowledgments:

The authors thank those involved in the development for their efforts, and recognize D. Tanksalava for help in setting the direction for the processor.

References:

[1] Hunt, D., "Advanced Performance Features of the 64-bit PA-8000," Compcon, Digest of Papers, 1995.

[2] Naas, B., "Memory Performance Features of the 64-bit PA-8000" Symposium Record, Hot Chips VII, 1995.

[3] Naffziger S., " A Subnanosecond 0.5mm 64b Adder Design," ISSCC Digest of Papers, pp. 364-365, Feb., 1996.

[4] Gaddis, N. et al., " A 56-Entry Instruction Reorder Buffer," ISSCC Digest of Technical Papers, pp. 214-215, Feb., 1996

[5] Colon-Bonet G., et al., " A Dual Floating Point Coprocessor with an FMAC Architecture," ISSCC Digest of Technical Papers, pp. 356-357, Feb., 1996.

0-7803-3136-2 / 96 / $5.00 / © IEEE

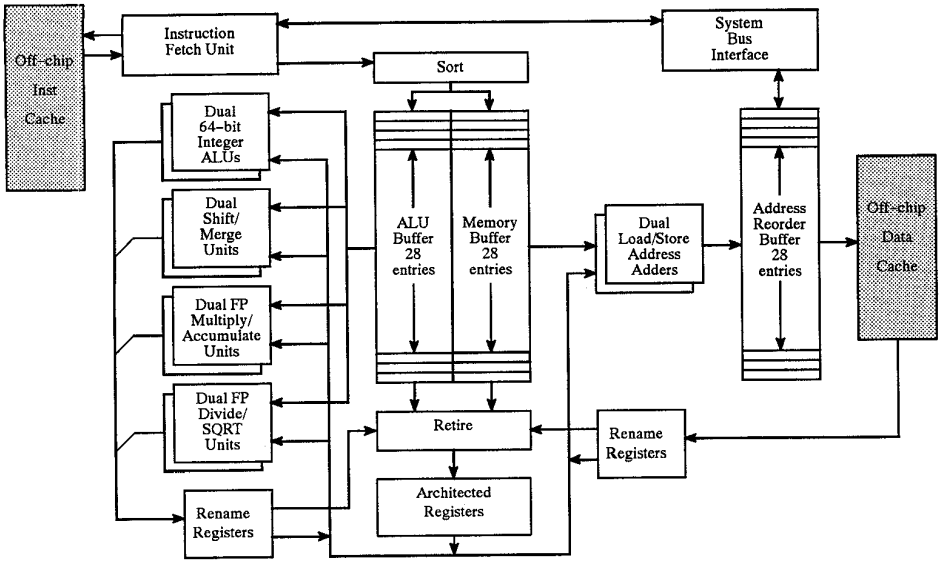

Figure 1: Chip block diagram.

Figure 2 and Table 2: See page 446.

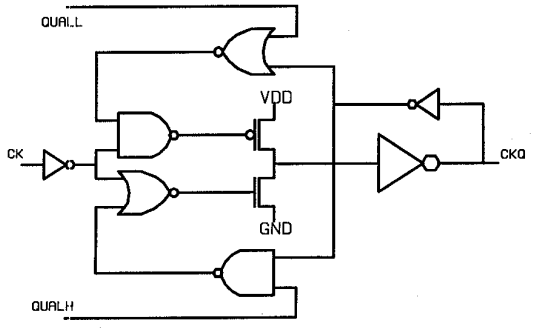

Figure 3: Clock gater circuit.

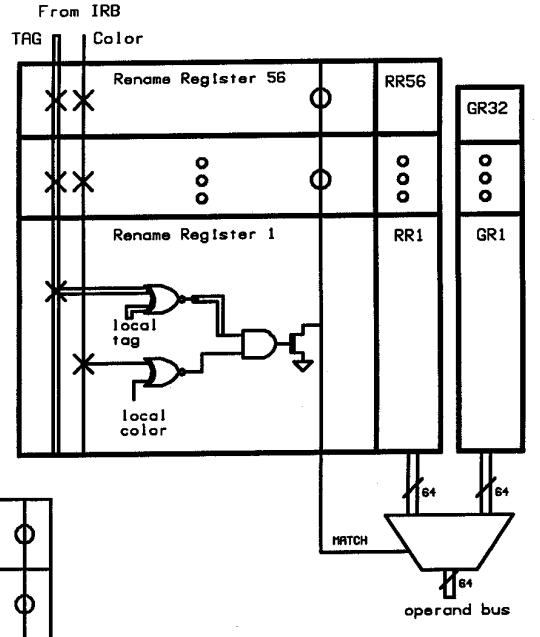

Figure 4: Rename register compare & operand data mux.

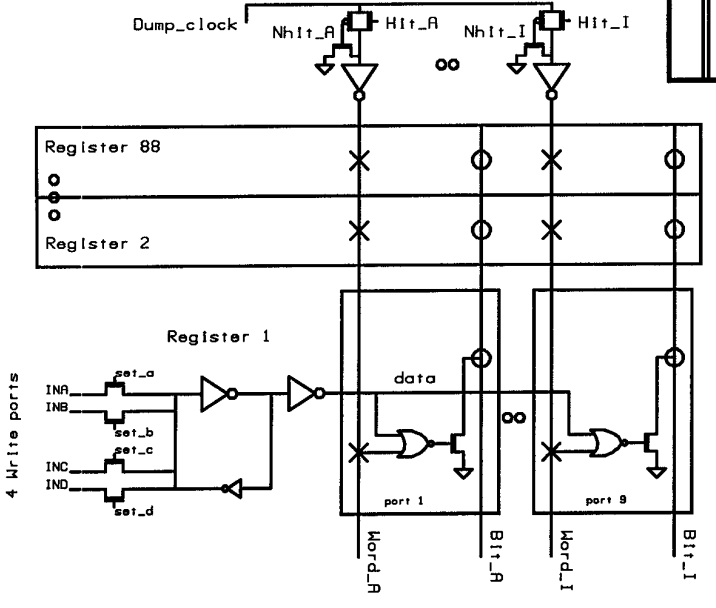

Figure 5: 4 write, 9 read port register.

Min. FET width, length	0.6µm
L effective	0.28µm
Metalization - 5-level	4 interconnect 1 power and clock
Metal 1, 2 width, space	0.6µm, 0.8µm
Metal 3 width, space	1.2µm
Metal 4 width, space	1.6µm
Metal 5 width, space	2.5µm

Table 1: Process specifications.

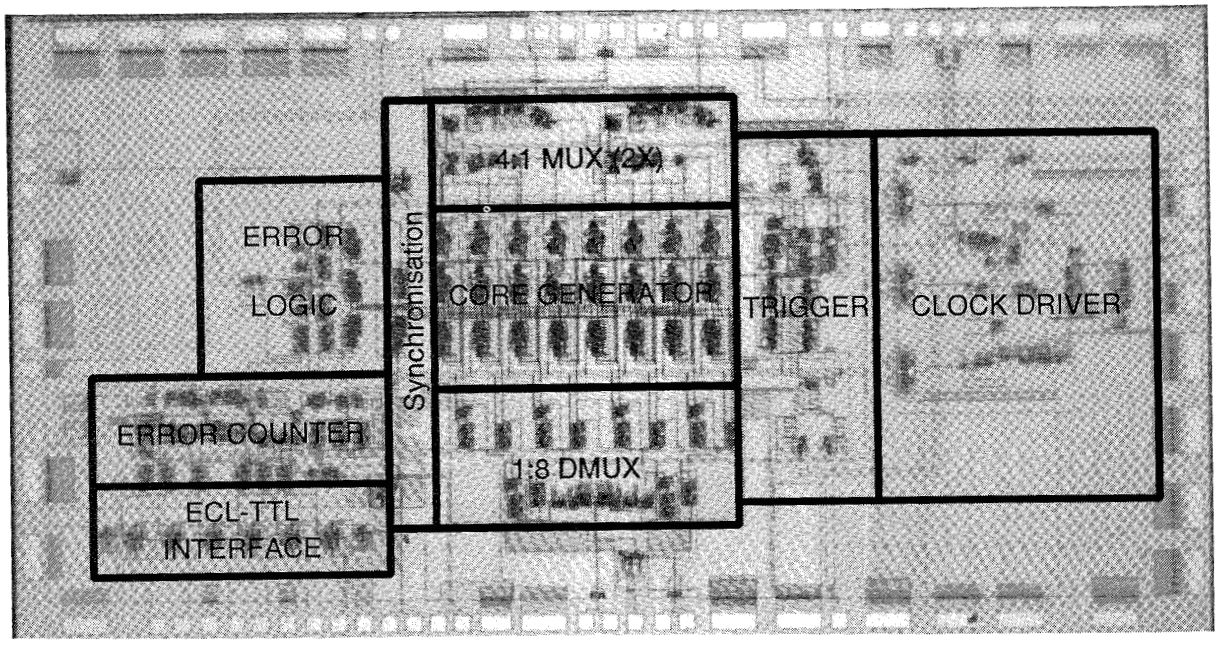

Figure 3: Chip micrograph.

FP 13.1: A Quad-Issue Out-of-Order RISC CPU

(Continued from page 211)

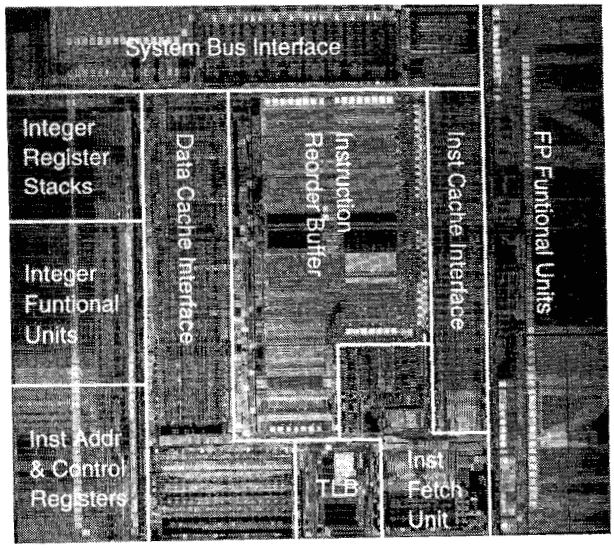

Figure 2: Quad-issue out-of-order RISC CPU micrograph.

Freq.(typical conditions)	150-250MHz
Pipeline	3-state instruction fetch
Dynamic reordering of instructions	(N-states)
	1 state execute (fully bypassed)
	2 state commit
Performance (estimated)	360 Specint92, 550 Specfp92
	8.6 Specint95, 15.0 Specfp95
Instruction reorder buffer	56 entry
Cache system	1MBx1MB split I/D external synchronous cache
	2-state latency (address to data)
Branch prediction cache	256 entry
Number of functional units	8
Peak sustainable throughput	4 instructions/cycle
Supply voltage	3.3V
Transistors	3.8M (no on-chip cache)
Die size	17.68x19.1mm^2
Pads	708 signals, 400 power/gnd
Package	flip chip/ceramic LGA

Table 2: Chip specifications.