Motorola Details Plan to Extend 68K Line

Superscalar 68060 is the Answer, But What is the Question?

By Linley Gwennap

At last month's Microprocessor Forum, Motorola countered Intel's scanty revelations of the Pentium (P5) chip with a detailed description of its own superscalar CISC processor, the 68060. The new Motorola chip, presented by Joe Circello, will use less power than Pentium and will probably be less expensive to build. For the first time, however, the 68060 will not be significantly faster on a clock-for-clock basis than its corresponding x86 counterpart. More importantly, if Motorola and Intel both met their schedules, the 68060 will be 12 to 18 months behind Pentium in reaching 66 MHz.

Circello said that the 68050 name will not be used because odd numbers indicate minor enhancements rather than major steps. Indeed, the 68060 represents a major step forward from the current 68040, particularly in the integer unit. As shown in Figure 1, the new chip will include dual integer execution units, allowing it to issue and execute up to two instructions per clock cycle with minimal restrictions. It uses a branch table and branch folding to reduce branch overhead.

The cache design is very similar to the 68040's (see µPR 2/21/90, p. 8). The separate on-chip caches are both doubled in size to 8K. The separate TLBs (Motorola calls them address translation caches, or ATCs) remain at 64 entries each, and both chips support a 32-bit virtual memory space. The caches are four-way associative and physically mapped. The data cache supports both writethrough and copyback modes.

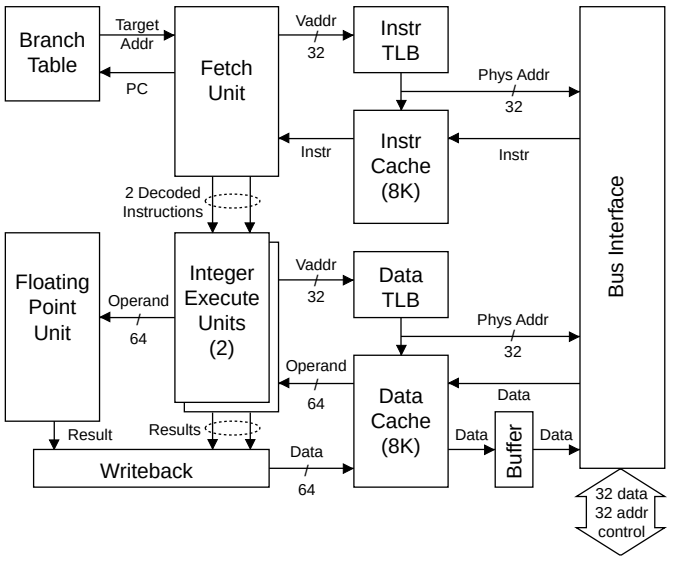

Figure 1. Simplified block diagram of the 68060. Some information has not been supplied by Motorola at this date.

There are a few changes from the 68040. The new data cache is four-way interleaved, or banked, allowing a read and a write to occur in the same cycle as long as they go to different banks; the 68040 cache stalls in this situation. The new design also includes a four-entry write buffer between the data cache and the system bus, which lets the processor generate a burst of writes to main memory without stalling the execution units.

The bus interface is a superset of the 68040 design (see µPR 3/21/90, p. 10); it adds a few new signals to simplify high-speed bus designs and removes the asynchronous (data latch enable) mode. It appears that the system bus is allowed to be clocked slower than the processor, although Motorola would not confirm this.

To integrate all of these features, over 2 million transistors will be required. The 68060 will use a 0.5-micron CMOS process with three layers of metal. The half-micron process reduces the operating voltage to 3.3V, although the chip will connect to standard 5V parts. Taking advantage of the lower power consumption of a 3.3V part, Motorola went further and added power management capabilities to the design. Power usage at full speed is expected to be 3–5 Watts, about the same as a 68040. Motorola claims the part will be offered at 50 MHz in 1H94, with a 66-MHz version to follow closely.

Instruction Pipeline

Figure 2 shows the 68060 pipeline, a stark contrast to the six-stage 68040 pipeline (see µPR 2/7/90, p. 1). The pipeline is broken into two parts: a semi-autonomous fetch unit and the execution units. The fetch unit reads instructions from the cache, partially decodes them, and places them into a large instruction buffer. The dispatch logic looks at the first two entries in the buffer, which is implemented as a FIFO, and issues one or both into the execution units. Because the fetch unit predicts most branches and continues to fetch instructions from the predicted path, the only interaction between the fetch and execute units occurs on mispredicted or non-predicted branches, which cause the FIFO to be dumped.

Most of the added stages are due to the complexities of decoding CISC instructions for superscalar issue. A new stage is added before the I-cache access to calculate the next PC; this calculation is more complicated due to the branch cache. Also, the 68040's "decode" stage has been broken into three parts. In "early decode," the

68060 partially decodes the available instructions to determine the general instruction type and register dependencies. During the next clock, this information is used to decide whether two instructions can be dispatched. Finally, an additional cycle is needed to fully decode the instruction and select operands for the AG stage.

All instructions are issued in order; the head of the FIFO must be issued into the primary execution pipe, while the next instruction can optionally be sent to the secondary pipe. The primary pipe is capable of handling any instruction, but the secondary pipe cannot execute floating-point instructions or branches. Because the data cache is single-ported, two instructions that both read (or write) to memory cannot be issued in the same cycle. Also, some instructions must be single-issued; for example, FP store instructions can cause late traps, and thus are not paired to maintain the precise exception model. Circello hinted that some instructions without memory references actually execute in the AG stage using the simple address-generation adder. Such an instruction could be paired with an instruction that depended on this result, increasing the number of paired instructions.

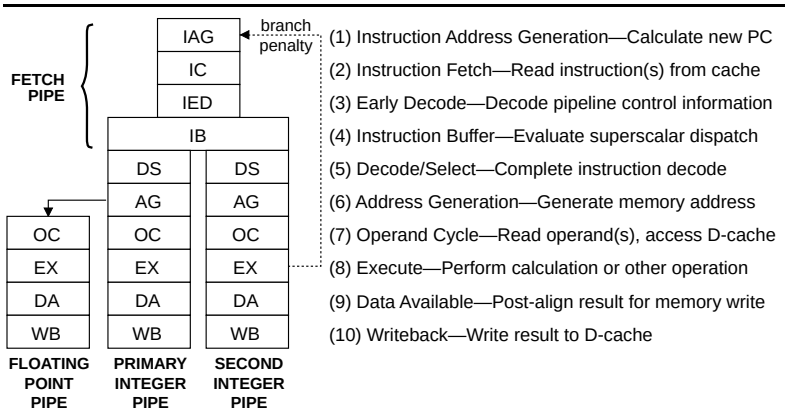

Figure 2. 68060 pipeline for fetch and execution units.

Motorola's goal is to provide high performance on existing binaries without recompilation. Based on traces of over three billion instructions, Circello expects the 68060 to issue 50%–60% of instructions in pairs. Figure 3 shows that issuing half the instructions in pairs results in a 33% performance increase due to superscalar issue; at 60%, the performance boost is 43%. These figures will be reduced due to cache misses and other overhead cycles, which are not improved by superscalar issue.

clock cycle:	1	2	3
PRIMARY	1	3	4
SECONDARY	2		

Figure 3. Pipeline utilization.

Branch Processing

One disadvantage of a long instruction pipeline is handling branches. In the 68060, the outcome of a conditional branch is not known until the end of the EX stage, resulting in the lengthy branch penalty shown in Figure 2. To avoid this penalty, Motorola added a 256-entry branch history table to the design. Unlike the branch target cache in the 88110, the 68060's branch table keeps track of recent branches by their instruction address, storing their actual target address with some history bits to indicate how often the branch has been taken. During the IC stage, the branch table is accessed in parallel with the instruction cache; if the table predicts that the current PC is a taken branch, the instructions being fetched from the I-cache are discarded and new instructions are fetched from the predicted PC on the next cycle. In this way, the branch is *folded* and does not occupy a pipeline slot, effectively executing in zero cycles as shown in Figure 4.

As in the 68040, non-taken branches are handled less optimally. If the branch table predicts that the branch is not taken, the branch remains in the instruction buffer and eventually takes up an instruction slot. The same thing happens if a branch instruction does not "hit" in the branch table, or for certain branches that are never predicted, such as subroutine returns.

On a conditional branch, the prediction mechanism works the same way, but the eventual resolution of the condition may invalidate the prediction. The processor keeps track of folded branches separately without using the normal pipeline; when the instruction before the folded branch completes the EX stage, the processor checks the condition code to verify the prediction. If the prediction turns out to be invalid, the processor must dump the instruction buffer and begin fetching from the correct location, causing the 7-cycle penalty shown in Figure 2.

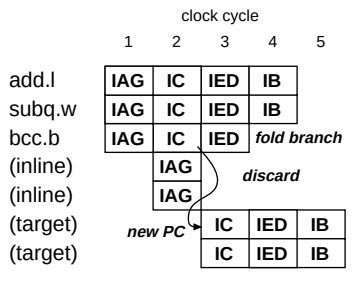

Figure 4. 68060 taken branch.

If the branch table results in about 85% correctly predicted branches, the average execution time of a taken branch would be about 1 cycle (7 cycles × 15% of branches) and 2 cycles for non-taken branches. This compares favorably with the 68040, which needs 2 cycles for taken branches and 3 cycles for non-taken.

Floating Point

Although 68060 floating point provides some improvement over the 68040, the new design does not include a state-of-the-art pipelined FP unit. Perhaps this reflects Motorola's lack of emphasis on the technical workstation market. Table 1 compares the throughput of various FP operations on the 68040 and 68060. The

INSTRUCTION (source, destination)		68060	68040	Pentium	i486
fmov.d	mem,reg	1	3	1	3
fmov.d	reg,mem	1	3	1	8
fadd.d	mem,reg	3	3/7	1/3	8-20
fmul.d	mem,reg	4	5/9	1/3	14
fdiv.d	mem,reg	24	38/42	38	73

Table 1. Floating-point latency (or throughput/latency for pipelined functions) in clock cycles for various CISC chips.

major improvement in loads and stores is due to a 64-bit path to the data cache. The 68040 holds a two-to-one advantage over the 80486 for most operations, but Intel's improvement on its Pentium processor should allow it to match the 68060.

One advantage that the 68060 will have over Pentium is that it can issue an integer instruction in parallel with most floating-point instructions. Also, it can continue to issue integer instructions into both pipes while a long-latency operation continues in the floating-point unit. Pentium, in contrast, ties up the entire processor when performing floating-point calculations.

The FPU is, of course, fully compatible with the 68040 and includes support for extended (80-bit) precision values. The cycle times in Table 1 are for double-precision values and may be longer for extended precision.

Pipeline Example

Figure 5 shows an example of instruction flow through the 68060 pipeline. This example shows the inner SAXPY loop from *matrix300*. Only the execution units are shown, since the fetch unit is decoupled, and the writeback stages have been deleted for simplicity. (Refer to Figure 2 for an explanation of the stage names.)

The first thing to notice is that the loop executes 13 instructions in 10 cycles, resulting in an IPC (instructions per cycle) of 1.3. Ten instructions are issued in pairs, two are issued singly, and the branch is folded; three cycles are spent waiting for the floating-point unit

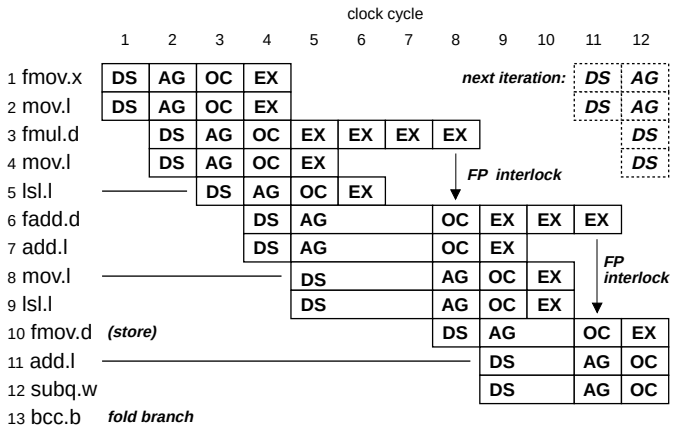

Figure 5. SAXPY loop executed in the 68060 pipeline.

Price and Availability

The 68060 is not yet announced; no pricing information is available. Motorola estimates that the chip will be available in 1H94, but this is not a committed date. A product brief (not a full data sheet) is available by ordering document MC68060/D. The *MC68060 User's Manual* is expected to be available by 3Q93; order number MC68060UM/AD.

Contact Motorola Literature Distribution at PO Box 20912, Arizona 85036; phone 602/994-6581, fax 602/994-6430.

in this FP-intensive loop. This single hand-picked example should not be considered typical, but it is real code generated by a 68040 compiler for a well-known program.

Instruction 3, an FP multiply, spends 4 cycles in the execute phase. While it is executing, however, instruction 4 is able to complete in the secondary pipeline. Furthermore, instruction 5, an integer instruction, is issued and completed without being hindered by the FP multiply, but instruction 6 cannot be issued into the secondary pipe because it is a floating-point operation. On cycle 4, the FIFO instruction buffer is advanced, and the top two instructions (6 and 7) can be issued simultaneously. The only other instruction that is executed singly is the FP store (10), which cannot be paired.

The long FP instructions freeze the pipeline during cycles 6 and 7, and also cycle 10. Notice that instructions that are further along than the frozen instruction can proceed to completion, but instructions in previous stages must wait until the freeze condition is cleared. Also, the freeze affects both integer pipelines.

Comparisons with Pentium

The block diagram of the 68060 is very similar to that of Pentium *(see **061405.PDF**)*. Both chips include two integer units, a floating-point unit, separate 8K instruction and data caches, and a branch history table. Although Intel has not specified Pentium's bus interface, they have stated it will be 64 bits, while the 68060 will use a 32-bit bus. On integer code, Pentium can fetch two 32-bit values from its on-board cache in a single cycle, while the 68060 will be restricted to a single read or write per cycle, but 68060 may have some limited issue of dependent instruction pairs. On floating-point code, Pentium's pipelined FP unit compensates for its inability to simultaneously issue integer and FP instructions. Thus, cycle for cycle, Pentium is likely to match the 68060 on a variety of benchmarks.

Pentium also has more aggressive clock frequency goals. Intel expects to ship Pentium at 66 MHz in 1Q93, and at 88 MHz in 1994. Motorola plans the first 68060s

to ship in 1H94 at 50 MHz, with a 66-MHz version to follow by the end of that year. Although either program could slip, Intel currently has plenty of working prototypes while Motorola has not yet tape-released the first 68060, giving Intel a significant lead.

In absolute terms, the 68060 could reach about 40 SPEC92 on both integer and floating point at 50 MHz, based on Motorola's goal of a 3.5× improvement over a 25-MHz 68040 without recompiling. Intel has been mum on performance, but we expect a 66-MHz Pentium to be about 30%–40% faster than a 50-MHz 68060. Some of Pentium's performance will be lost without recompiling, however, particularly for floating-point code.

One advantage Motorola may have is that its chip will use a third fewer transistors than Pentium and will use a half-micron CMOS process, compared to Intel's 0.8-micron BiCMOS process. As a result, the 68060 should be smaller and less expensive to manufacture than Pentium, even when Pentium moves to a half-micron process.

Who Will Buy?

Although it is interesting to compare the two leading CISC designs, in the PC world, the choice comes down to Macintosh versus Windows or DOS, and the microprocessor doesn't really matter. Apple's plans to move to PowerPC systems could put a damper on sales of future 68060 Macintoshes, although the common bus interface with the 68040 will make it easy for Apple (or vendors of accelerator boards) to offer such a product for those who need binary compatibility.

In the workstation market, its competitors are RISC chips. After Motorola's debacle in getting the 68040 into production, this competition is already lost; the only remaining major vendor of 68K workstations, HP/Apollo, has said that all of its new systems will use PA-RISC. Any vendors thinking of sticking with the 68060 will notice that today's leading RISC chips offer nearly twice the integer performance and over three times the floating-point performance of the 68060 more than a year before it is expected to be available. One general-purpose market where the 68060 may be popular is the VME board market.

Motorola has thoughtfully added features that make the 68060 more attractive for portable systems. Low-voltage operation will reduce power consumption. To further reduce power, the control logic will disable functional units that are not being used on a cycle-by-cycle basis. The design will be fully static, allowing the clock to be reduced or stopped. A new instruction, LPSTOP, facilitates power management by halting all internal processing until an interrupt or reset signal is received. The 68060 will still be too power-hungry and too expensive for handheld systems, but it will perform well in PowerBook-type systems. The PowerPC 603, also due in 1H94, will compete for these types of products.

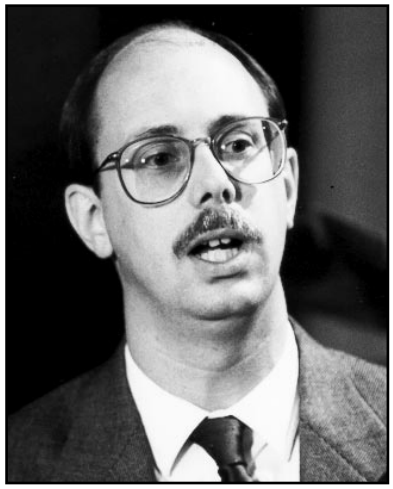

PHOTO BY ANNE HAMERSKY

"There was a very deliberate set of decisions made in a lot of respects to try to minimize the complication with this particular instruction set in terms of the superscalar approach ….We were willing to give up whatever incremental performance we might have been able to get through things that were a lot more complex, in the goal of simplifying the design."

Joe Circello, Motorola

The 68K family continues to have success in the embedded market, which will be the best target for the 68060. The new chip is designed so that it can be easily stripped down for use in the embedded market. The floating-point unit and MMU can be removed, and the cache sizes can be modified. As with previous generations, we can expect to see "EC" versions for embed-ded control that cost half as much as the full-fledged 68060. Eventually, the integer core will be added to the 68300 line, which today combines 68000 and 68020-like cores with a set of integrated I/O devices.

Most embedded customers today are using 68000s and 68020s; it will be years before these customers need the firepower of the 68060. To look at it another way, it will require at least one process shrink before the 68060 is affordable for these price-sensitive customers.

There are always speed addicts who need the newest CPU for their network gateways, X-terminals, and high-end page printers. Since these vendors don't have to support an installed base of customer applications, many are evaluating the 68K products against RISC chips such as Intel's i960, AMD's 29K line, and embedded MIPS variants. Motorola will have to aggressively price its 68060 product family to compete in this environment.

Motorola's sad experience with the 68040 shows how fast a company can lose a market that it once dominated. The high-performance workstation market has moved entirely to RISC, and Apple is following suit. For these vendors, the 68060 is a stairway to nowhere. Motorola retained its more price-conscious embedded customers, most of which have not even started designs with the 68040. For these customers, the 68060 demonstrates that there is a long-term growth path for the 68K. ♦

