

SGI Provides Overview of TFP CPU

New Chip Set Boosts Floating-Point, Used in Parallel Supercomputer

By Linley Gwennap

Silicon Graphics, while announcing a powerful new multiprocessor system, revealed the first details of its forthcoming TFP processor. The new chip set, which SGI expects to ship around the end of the year, is a four-way superscalar design. It can dispatch any combination of instructions among an impressive array of two integer units, two floating-point units, an integer multiplier, two load/store units, and a branch unit. TFP is the first superscalar implementation of the MIPS architecture.

The processor is divided between two chips, as shown in Figure 1. The floating-point chip, rumored to use a Weitek design, includes the FP registers and two complete floating-point units (FPUs). The integer CPU chip contains the instruction dispatch unit, the remainder of the functional units, 16K of instruction cache, and 16K of integer-only data cache.

The cache design is unusual in that the first-level data cache, located on the integer chip, does not cache floating-point data; this is because the FP registers are located on the other chip. All FP data is kept in the off-chip cache, which can be 2M–16M in size. This cache, which is four-way set associative, also serves as second-level storage for integer data and instructions. A 128-bit

bus provides peak transfer rates of 1.2 Gbyte/s between this “streaming” cache and the on-chip data storage. Two-way interleaving helps the external cache achieve this bandwidth. The cache’s large size and high bandwidth to memory should give TFP an advantage in applications with large data sets.

To keep the instruction unit fed, branch prediction is used to execute past branches by fetching instructions from the most likely path. SGI would not discuss any details of the prediction technique, or whether the processor supports speculative execution.

The complex superscalar capabilities—particularly the two FPUs—forced the designers to use four million transistors, more than could be placed on a single chip even with 0.6-micron CMOS technology. The resulting two-chip design created interchip communication delays that helped keep the expected clock rate to just 75 MHz, half that of the R4400. The complicated instruction dispatch may also curtail the clock rate. Even at 75 MHz, TFP would still be faster than any similarly-complex design, such as IBM’s RIOS (62.5 MHz), SuperSPARC (40 MHz), or the 88110 (50 MHz).

SGI indicated that TFP chips will be made available to other system vendors at about the same time as they are used in the company’s own systems. Toshiba will fabricate the parts for SGI, and the design will be offered to the other MIPS foundries as well. Samples should be available in the fall, and volume shipments are planned for 1Q94.

TFP implements a new instruction set, MIPS-IV, that features extensions to the R4000 architecture. These include a multiply-add instruction (similar to PA-RISC and POWER), conditional move instructions (similar to Alpha and SPARC-V9), and a “register+register” addressing mode. TFP will, of course, be backward-compatible with all user code compiled for previous MIPS processors.

The centerpiece of the new processor is the dual floating-point units, each of which is capable of starting a multiply-add combination every cycle. At 75 MHz, this provides a theoretical peak rate of 300 MFLOPS. SGI would not provide any benchmark data, but SPECfp92 performance of 200–250 is feasible with this design. It is clear that the new processor will more than double the FP performance of the 150-MHz R4400.

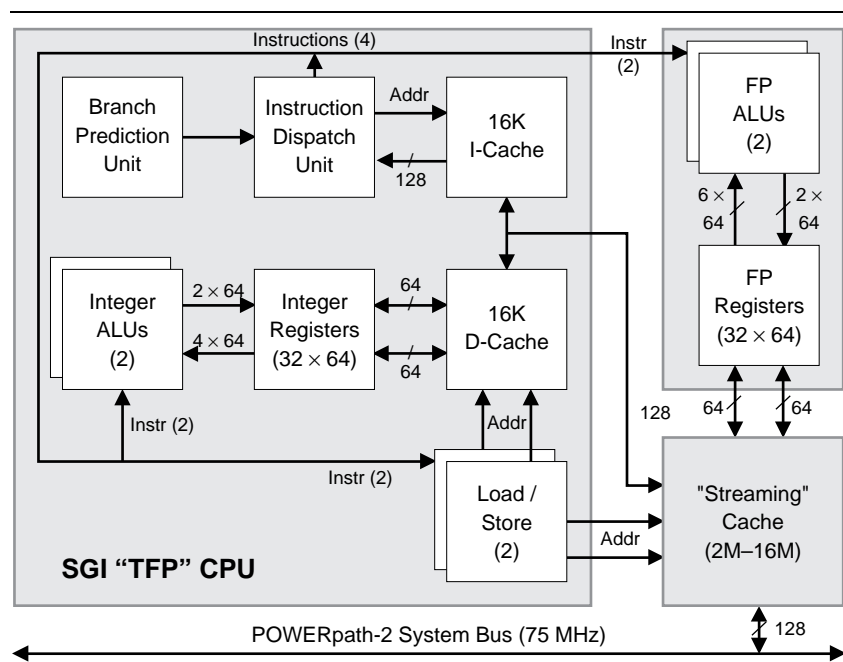


Figure 1. Block diagram of TFP processor. The CPU chip is at left, with the FP chip in the upper right. The cache uses multiple SRAM chips.

Strong Competition Ahead for TFP

Whether TFP will outperform other vendors' chips remains to be seen. For pure performance, the toughest competition may come from IBM's RIOS-2 processor (*see 0702MSB.PDF*). Like TFP, the IBM design also features a large set of functional units and dual FPUs that can simultaneously execute multiply-and-accumulate instructions. By using MCM packaging, RIOS-2 will probably exceed TFP in clock frequency and could beat its performance as well. IBM is already testing first silicon for RIOS-2, while SGI is nearing tape release for TFP, so the POWER processor is likely to reach the market first. RIOS-2, however, is not expected to be sold on the merchant market.

HP and DEC, the current floating-point leaders, are not likely to give up that lead easily. Today's PA7100 and Alpha processors are both capable of nearly 200 MFLOPS peak. DEC already has a 300-MHz Alpha planned with a similar release date as TFP, and HP will certainly deliver improved performance by then. By sticking with a single-chip design, DEC and HP should have a significant cost advantage over TFP (and RIOS-2, for that matter).

TFP's cost disadvantage will be exacerbated by low volumes, since the new chip will not replace the R4400 as the flagship MIPS processor. Although TFP will offer superior floating point, the integer performance of the two chips will be roughly the same, assuming that TFP's superscalar design allows for an average of twice as many instructions per cycle as the R4400. Thus, TFP will be restricted to high-end applications that require its immense floating-point power.

New Parallel System Features TFP

SGI's first TFP-based product will be the highly-parallel Power Challenge XL supercomputer. The system supports up to nine processor boards, each of which contains two complete TFP processors, in a standard 1.8-meter cabinet. Thus, the maximum configuration supports 18 TFPs and is rated at 5.4 GFLOPS of peak performance. A deskside system, the Power Challenge L, supports up to three processor boards.

The same packages also support processor boards with four R4400 processors each; in this case, the systems are labeled "Challenge" without the Power. The Challenge XL can hold up to 36 R4400 processors. Because of the lower floating-point performance of the R4400, the company is marketing these systems as high-end servers rather than supercomputers. Although SGI focuses primarily on the technical market, these new servers compare favorably to commercial database engines such as Sequent's 32-processor Symmetry system and Sun's new 20-processor SPARCcenter 2000.

Like most smaller parallel systems, the new

Challenge line implements a shared-memory model to simplify the programming task. The XL systems support up to 16G of real memory that can be interleaved eight ways to increase its bandwidth, while the L versions allow up to 6G. The processors, memory, and I/O are connected through the 1.2-Gbyte/s "POWERpath-2" system bus. The I/O system allows for standard serial, parallel, and Ethernet connections as well as higher-speed devices such as SCSI-2, VME64, FDDI, and HiPPI.

These systems run the IRIX operating system, the same OS used by SGI's workstation products. This UNIX-derivative operating system has been enhanced with a multi-threaded kernel and parallelizing compiler technology to simplify the creation of parallel programs.

The Power Challenge systems offer performance equivalent to current supercomputers, but the price is much lower; an entry-level configuration of the L version, with two TFP processors, has a list price of under \$120,000. These Power systems are expected to begin shipping in early 1994. The R4400-based Challenge systems will ship later this quarter and are upgradeable to TFP processors as soon as TFP becomes available.

IBM's recently-announced POWERparallel system will provide a significant challenge to SGI's Power line. (Although the similar names will undoubtedly confuse customers, SGI came up with the "Power" name first.) IBM's system couples up to 64 RIOS-1 processors and is expected to ship this fall. Although the entry price of about \$300,000 is not as low as SGI's, the systems are comparable in dollars per MFLOP. IBM will surely offer its RIOS-2 processors in the same package at some point, creating a truly formidable system.

Conclusion

The TFP processor will bring the MIPS architecture out of the doldrums in floating-point performance and let it sail among the leaders. By announcing the chip a year before shipments, however, SGI risks taking its best shot and then watching other processor vendors meet or beat TFP's performance. The expected high price of the new processor likely will relegate it to a small niche in scientific computing.

The new processor is important for providing SGI with an entry into the supercomputing arena. By leveraging its RISC technology and its focus on the technical market, the company should do well. IBM and other RISC vendors have similar ideas, however, and will be strong competitors, along with established vendors of massively-parallel systems. Advanced compilers and operating systems will be the key that unlocks the power of these highly-parallel systems; SGI has some solid technology in these areas that will boost its chances for success. Traditional supercomputer and minisuper vendors such as Cray, NEC, and Convex are in for some trouble. ♦