Pentium PC Performance Varies Widely

Fast Cache, Memory Needed for Good Performance on SPEC

by Linley Gwennap

Like most processor vendors, Intel usually touts a single SPECmark number for its Pentium processor, one obtained from a high-end server system. A PC motherboard, however, will not match this level of performance. In fact, the PC performance varies significantly, depending on the size and speed of the cache and memory systems. Recently released data from Intel shows that a high-end PC comes close to matching the SPEC performance of Intel's server motherboard. Low-cost PCs, in contrast, may be 20–30% slower (or more).

Good PC Motherboard 10–15% Slower

Figure 1 compares the performance of three Intel motherboards running the same SPEC binaries. The highest performance is obtained from the Xtended Xpress board (XXpress) with 1M of two-way associative secondary cache, which Intel uses for its published SPEC ratings. Slightly lower performance is obtained when the same board is configured with 512K of associative cache. The XXpress, designed for servers, uses Intel's 82496 cache controller and 82491 cache RAM to provide zero-wait-state (2-1-1-1) cache accesses for the most recent set and 3-1-1-1 accesses for the second set.

The third board, a standard Intel motherboard aimed at desktop PCs, contains the Triton chip set and a 256K secondary cache. This cache uses synchronous SRAMs to deliver 3-1-1-1 performance, only slightly slower than the cache on the XXpress board. Triton supports a direct-mapped cache; the smaller size combined with the lack of associativity gives the Triton cache a much lower hit rate.

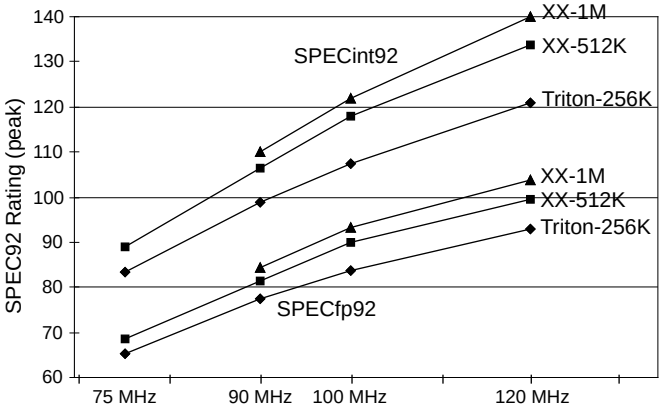

Figure 1. With a 90-MHz Pentium, the performance of Intel's standard Triton motherboard is within 10% of that of the Xtended Xpress (XX) motherboard with either 512K or 1M of secondary cache. The gap widens, however, with a 120-MHz Pentium. (Source: Intel)

There are some differences in the main-memory configurations as well. The XXpress uses interleaved fast-page-mode DRAM to achieve a 14-3-3-3 access to main memory. Using EDO DRAM, the Triton board accesses main memory at a 12-2-2-2 rate. Thus, the Triton board actually has a lower latency and higher bandwidth to main memory than the XXpress.

With a 90-MHz Pentium, the Triton board delivers about 10% less performance on SPECint92 than the XXpress board with 1M of cache. Most of this difference is due to the higher cache miss rate. As the CPU speed increases, the penalty of each L2 cache miss becomes proportionally greater. This factor increases the performance gap to 14% when the CPU moves to 120 MHz.

On SPECfp92, the gap is 8% for a 90-MHz system, increasing to about 11% with a 120-MHz Pentium. But while the integer performance increases almost linearly with clock speed, particularly with 1M of L2 cache, the SPECfp92 curves start to flatten out as the clock frequency increases. Because SPECfp92 benchmarks have a relatively high cache-miss rate, even with a 1M cache, their performance becomes limited by the bus bandwidth rather than the CPU speed.

The biggest change is seen between 100 and 120 MHz. The 100-MHz Pentium uses a 66-MHz bus, while the 120-MHz part is limited to a 60-MHz bus. Thus, the slower CPU paradoxically has 10% greater bandwidth. As noted, the bus bandwidth has a significant impact on SPECfp92 performance.

The Triton board handles up to 512K of secondary cache. This configuration, not shown in the figure, adds 2–4% performance on SPECint92 compared with a 256K cache and is more helpful at higher clock speeds. With a larger cache, the PC motherboard comes within 8–10% of the performance of the fastest XXpress.

Low-Cost Memory Reduces Performance

Pentium PC performance is reduced with non-EDO (fast-page-mode) memory. Compared with the Triton board in Figure 1, this lower-cost option erases 3–6% of integer performance, with the biggest impact coming at the highest clock speeds. The impact is smaller on systems with a 512K cache, as they have fewer accesses to main memory.

Intel did not provide SPEC ratings on any systems with asynchronous L2 caches. A typical asynchronous cache returns data at a 3-2-2-2 rate, taking nearly twice

as long to refill a single cache line as a synchronous cache would. The asynchronous cache, popular in low-cost systems, causes a performance drop of about 10% or so. As with other changes, the impact is biggest on the fastest processors. It is also worse when combined with fast-page-mode DRAMs than with EDO.

A few Pentium systems go entirely without an L2 cache; such a system could take a 20–30% performance hit compared with the Triton system in Figure 1. As always, the impact is worse for faster Pentiums, which is why cacheless Pentium systems nearly always run at 75 or 90 MHz. With EDO DRAM, the impact is reduced significantly, because the fast memory helps compensate for the lack of secondary cache. Intel reports that, with EDO, the performance loss is less than 20%.

As both synchronous SRAMs and EDO DRAMs are rapidly increasing in volume, their prices are dropping. Intel believes that, once these prices reach maturity, the most popular low-end configuration will be a cacheless system with EDO memory. Midrange systems will use a 256K synchronous cache with fast-page-mode or EDO DRAM, while high-end systems will use 512K synchronous caches and EDO DRAM.

High-End PC Does Well on SPEC95

Intel has reported more limited testing with the new SPEC95 benchmark suite *(see **091102.PDF**)*. Table 1 compares baseline SPEC95 scores for a Triton motherboard with a 512K synchronous cache and EDO DRAM against the XXpress with 1M of cache. Both rely on a 133-MHz Pentium CPU.

On the new integer suite, this Triton board comes within 8% of the XXpress, roughly the same as on the SPECint92 suite. On SPECfp95, the 6% gap is also similar to SPEC92. On a few individual tests, Triton even comes out ahead, due to its better memory bandwidth.

The SPEC95 suite generates more cache misses than SPEC92, making it more representative of PC application performance. We believe this effect will reduce performance on Triton boards with less expensive memory configurations, such as a smaller cache, asynchronous SRAMs, or fast-page-mode DRAM. Intel has not provided SPEC95 measurements in these configurations. Thus, the SPEC92 results described above may slightly underestimate the performance gap between a lower-cost Triton board and a maxed-out XXpress when running typical PC applications. Note that neither SPEC suite tests graphics or disk performance.

As Always, *Caveat Emptor*

These figures show that performance comparisons are a tricky business. Intel likes to compare the SPEC scores achieved on its Xtended Xpress board against RISC workstation scores. These numbers are not accurate if you want to compare a high-end Pentium PC against a RISC workstation, because a Pentium PC with a Triton motherboard delivers lower performance than the XXpress board.

For a true high-end PC with 512K of synchronous cache and plenty of EDO DRAM, the performance difference is less than 10%. Even with this degradation, the PC can outrun many workstations that cost $10,000 or more. But the PC's performance drops rapidly with a more mainstream configuration, such as a 256K asynchronous cache and fast-page-mode DRAM. Chip sets other than Triton may reduce performance further, due to slower memory or PCI performance.

Applying a single performance number to a microprocessor is always dangerous. A maximum "does not exceed" rating can be useful in comparing one processor to another, but this performance can easily be squandered by a slow system design. SPEC ratings for specific RISC systems are generally available. Few PCs, however, have published SPEC ratings. The guidelines described above can help PC buyers estimate the performance of a Pentium system instead of simply using the maximum SPEC ratings. ♦

For More Information

For more information on Intel's Pentium processor performance, check the Web at *www.intel.com/procs* or contact your local Intel sales office.

	Triton w/ 512K L2	XXpress w/ 1M L2	Percent Change
099.go	3.98	4.62	+16%
124.m88Ksim	2.86	3.04	+6%
126.gcc	2.96	3.34	+12%
129.compress	3.46	3.63	+4%
130.li	3.92	4.23	+7%
132.ijpeg	2.89	2.86	-2%
134.perl	3.74	4.03	+7%
147.vortex	3.01	3.37	+11%
SPECint95 (base)	3.32	3.60	+8%
101.tomcatv	3.18	3.40	+6%
102.swim	3.50	4.08	+16%
103.su2cor	1.38	1.65	+19%
104.hydro2d	1.71	1.60	-7%
107.mgrid	1.29	1.29	0
110.applu	1.25	1.23	-2%
125.turb3d	2.65	2.71	+2%
141.apsi	2.76	2.97	+7%
145.fpppp	4.52	5.28	+16%
146.wave5	3.22	3.62	+12%
SPECfp95 (base)	2.32	2.48	+6%

Table 1. A Dell system with a Triton motherboard and 512K of synchronous L2 cache achieves about 8% less integer performance than an Xtended Xpress motherboard with 1M of cache, when comparing SPECint95 (baseline) scores. On SPECfp95, the gap is even narrower. (Source: Intel, SPEC)

