Gigahertz processors compared
HIGH CLOCK RATES

Processors with clock speeds of 1000MHz were announced months ago but have only recently begun appearing in stores. We investigate the effects of the higher clock speeds under Linux. Bernhard Kuhn gets off the starting blocks.

New processors with high clock frequencies seldom work satisfactorily with older motherboards. Thus in addition to the Coppermine processor (Slot 1 with 256 KByte on-chip cache), Intel loaned us a VC820 motherboard with an audio/modem riser slot. Also bundled was a high-grade 128 MByte RDRAM memory module (Rambus DRAM) from Kingston Technologies. A dummy is required for the unused second RIMM slot (Rambus Inline Memory Module) to ensure the whole thing worked – unlike SDRAM, Rambus is extremely sensitive to such things (see Fig. 1).

AMD loaned us an entire computer for test purposes. They too have not yet achieved the proclaimed socket changeover – the processor on our test machine was an old K7-based chip with 512 KByte of external level 2 cache on the slot A motherboard. AMD provided 384 MByte of main memory distributed over three slots, which we reduced to that of the Intel system in order to achieve fair conditions for testing. The remaining components on the Irongate chip set AMD motherboard did not have any significant effect on the test results.

With floating-point-intensive benchmarks such as the "Blenchmark2" (Figure 2, Blender Version 1.80a) and the Povray-Skyvase-Test (Figure 4), the Athlon chalked up a clear lead: 15 to 20% better performance compared to Intel. This is fairly remarkable.

With the "Developer Benchmark" the result is the other way round. The kernel compilation (see box and Figure 4) was completed by the Intel chip roughly 20% faster than from the AMD. A similar response was seen with Ralph Hlsenbusch's performance test from the "Shared Services Benchmark Association" (iX-SSBA), version 1.21E. However, in this combined server/application test with a runtime of about ten minutes (564 seconds), a lead of 18 seconds is poor testimonial for the RDRAM, which is supposed to be ideally suited to these tasks. With the *stream* benchmark we again found no evidence of the alleged two to three times higher transfer rate of RAMBUS compared to conventional SDRAM, and *nbench* verified AMD's position as a wonder weapon for being in the black even in the higher memory index (6.0 versus 4.5 with Intel).

Duel with the gladiator

An Elsa Gladiac (32 MByte DDR-RAM) was used for the graphics-dependent performance tests with 24 bit colour depth. The combined force of Nvidia's Geforce 2 GTS chip and Intel's Pentium III allowed the highest 2D speed index recorded so far (see Figure 5). The Quake3 results in Figure 6 and the SPECviewperf-Suite cannot be compared directly, however. Unfortunately, the Intel equipment was loaned to our lab for only one week and our tests were performed with the then current SGI/Nvidia driver 0.93. A week later when we came to test the AMD chip it proved not to like this driver, so an update was necessary – forcing us to use the better performing version 0.94. As Figure 7 shows, this had scarcely any effect at all for high resolutions: in the Q3 demo even a Pentium III 550 MHz can easily keep pace with the gigahertz systems we had on test.

Fig. 1: To avoid trouble, Rambus requires all the RIMM slots to be filled, using a dummy if necessary

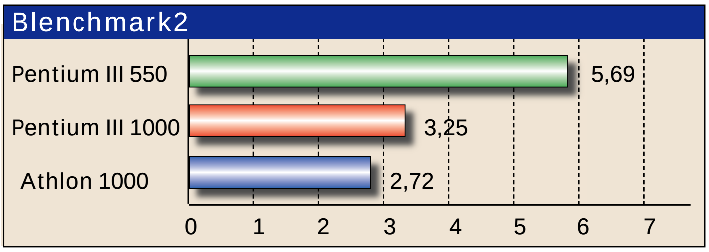

Fig. 2: The Athlon's floating-point units are just in the lead.

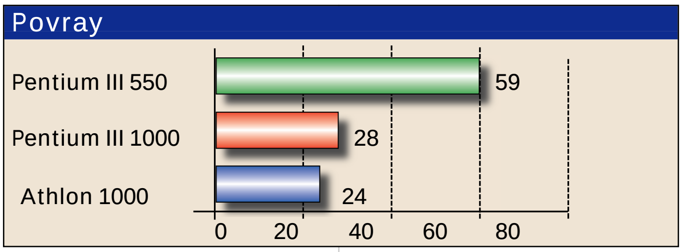

Fig. 3: AMD's Athlon is almost as quick as two 667MHz Alpha 21264 processors combined (22 seconds for pvmpov).

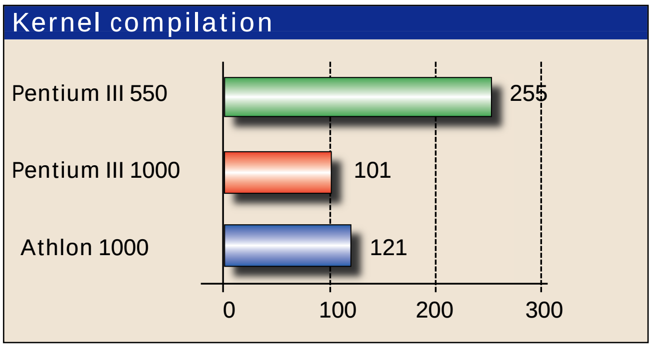

Fig. 4: For Linux developers Intel's Pentium III is only slightly more suitable.

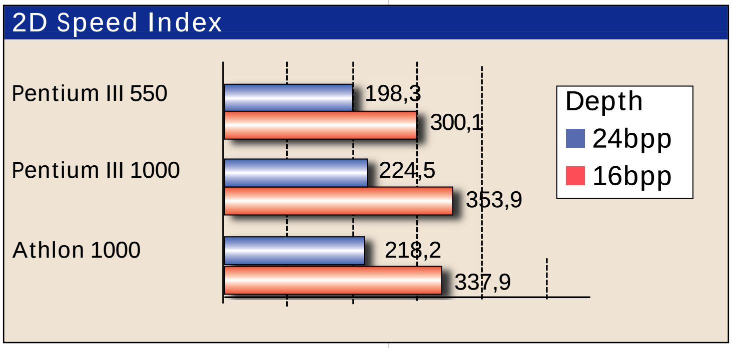

Fig. 5: An amazing combination: Pentium III and Elsa Gladiac

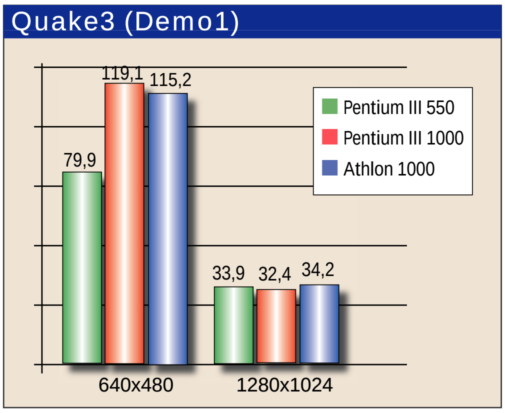

Fig. 6: During the textures computation even a weaker processor can keep pace with the four pixel pipelines of the Nvidia flagship.

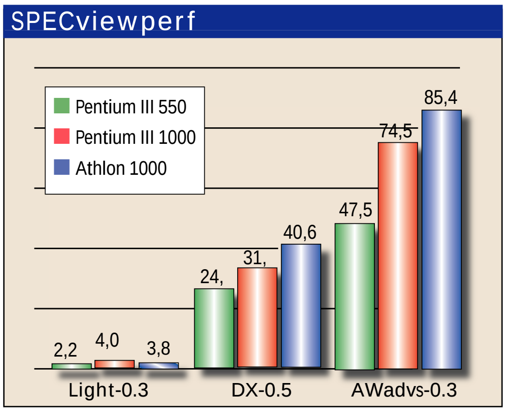

Fig. 7: Unfortunately these figures cannot be compared directly: in the AMD system a newer version of the graphics driver had to be used

Incidentally, for computers with AMD processors, the beta status of the SGI/Nvidia driver frequently made itself apparent with X-server freezes or entire system crashes. It wasn't possible to reproduce this behaviour with the gigahertz Pentium, suggesting that SGI and Nvidia possibly develop and test their Linux drivers with Intel systems only.

To sum up

As was to be expected, both the gigahertz systems are faster than their predecessors. But they were also significantly more expensive. Whilst AMD has an enormous lead over its competitors, there would really need to be compelling reasons to pay such a price – more than three times higher – for a 20% increase in performance compared to a machine running at 800 MHz. On top of this you must consider the cost of RIMMs for Intel's RAM bus. And applications that fully exploit this type of memory are few and far between.

In the end, fast processors are like fast cars. Common sense by and large compels most of us to choose a medium-class product even though we know they will provide less pleasure at the wheel. ■

Sense and nonsense about the kernel compile benchmark

USENET newsgroups frequently refer to the kernel compilation time as a measure of a computer's performance. Whilst the result reflects perhaps the suitability of a system as a development machine, without specifying the exact configuration along with compilation time, the times people come up with are meaningless as a basis for comparison. Depending on the kernel version, the number of features compiled into it and also the compiler options (and version) used, the time to complete the task can vary enormously. Consequently for the "Developer Benchmark" we used a virgin kernel 2.2.0 with the default configuration. Following a make menuconfig; make dep the time for the compilation is derived from the arithmetic mean of the execution times (elapsed) of three passes of the command time make bzImage.

■