

Intel Clock-Doubler 486 Debuts as 486DX2

486DX2 Aimed Primarily at OEMs — End-User Version to Follow

By Michael Slater

Intel has formally introduced the first version of its 486 microprocessor with an on-chip clock doubler. The chip operates from a 25-MHz external clock and appears nearly identical to a 25-MHz chip from a hardware perspective, but it operates internally at 50 MHz as long as its memory needs are satisfied by the on-chip cache. Whenever a cache miss, an I/O access, or a write cycle with the write buffer full occurs, the chip slows down to the external clock rate.

Intel has demonstrated clock-doubler 486 chips for more than six months, so the only real news is the name (486DX2), the price (\$550 in thousands), and the availability (now, for the 25/50-MHz version). The price is \$60 less than the standard 50-MHz 486 (priced at \$610), but \$133 more than a 33-MHz 486 (priced at \$417). The 486DX2-50 will be joined in the second half of this year by the 486DX2-66, which plugs into a 33-MHz 486DX socket and runs at 66 MHz internally. Intel has not committed to offering a 486DX2-100, which would use a 50-MHz external clock.

The 486DX2, previously known as the P24, is aimed primarily at OEMs, not at the upgrade market. It provides system makers with a very easy way to introduce a new model simply by replacing the 486DX in existing 25-MHz systems with a 486DX2. These "pseudo-50-MHz" systems are likely to displace the 486DX-33 as the most popular system for power users. True 50-MHz systems have been too expensive to become mainstream products and have been popular primarily as servers.

In the past, PC makers have offered "one of everything," but the increasing variety of 386/486 processors will require that they become more selective. Many vendors may choose not to offer true 50-MHz systems. Some might drop their 33-MHz systems, hoping to move users up to the 486DX2-50, but the current \$133 price differential between the two processors means that the difference in system cost is likely to be about

\$300—enough to justify including both models. The 486DX2-50 systems should be significantly less expensive than true 50-MHz systems, even though the processor itself is only \$60 cheaper, because slower cache memories and other system components can be used. The design task is also much easier; while designing a true 50-MHz system is difficult (unless Intel's expensive processor/cache module is used), the 486DX2-50 is a trivial upgrade to a 25-MHz design. This will enable every clone vendor to offer this configuration, while many of them have not offered 50-MHz systems, so pricing will be more competitive.

New systems using the 486DX2-50 should appear very quickly, since almost no system redesign is involved and Intel claims to have the chips ready for immediate delivery in volume; at least 15 system vendors are expected to announce systems simultaneously with the chip's announcement. Assuming Intel does not have major yield problems at the higher clock rates, 486DX2-66 systems should appear by fall Comdex.

The 486DX2 is fabricated using the same 0.8-micron, three-level-metal process as the 50-MHz 486. This is Intel's most advanced process, and the cost of fabricating chips on it is considerably higher than on the older 1-micron, two-level process used for the 33-MHz and slower 486 chips. The die size of the 486DX2 is identical to that of the 486DX-50; while it includes a small amount of additional logic to handle the half-speed bus interface, Intel's designers were able to fit it in the same area.

Performance

The 486DX2-50 is pin-compatible with the 486DX-25, and its signal timing specifications are identical. Optimum instruction times (assuming cache hits for both instruction and data accesses) are, of course, half that of the standard 25-MHz 486. The actual performance gain seen by the user is highly application-dependent, and can vary from as little as 10% for "cache-buster" programs or programs that are I/O-bound, to nearly 100% for programs that spend most of their time in short loops operating on small data sets.

Table 1 shows Intel's benchmark data for the DX2-50 and DX2-66, compared with the DX-33 and the DX-50; Figure 1 shows the same data graphically, normalized to DX-50 performance. The DX2-50 comes surprisingly close to DX-50 performance; it is less than 2% slower on the Norton SI and Dhrystone benchmarks, which fit in the on-chip cache. On the SPEC

	486DX-33	486DX2-50	486DX-50	486DX2-66
Norton SI V5.0	72.1	106.1	108.1	144.2
Dhrystone 2.1	43.5	64.0	65.4	85.5
SPECint89	19.0	25.4	28.5	34.0
SPECfp89	12.2	15.9	18.3	21.2
Linpack DP MFLOPS	1.6	2.3	2.4	3.1

Table 1. Benchmark results for 486-family processors.




Figure 1. Performance for 486 family members, normalized to 486DX-50.

integer benchmarks, the DX2-50 is 11% slower than the DX-50, and on the SPEC floating-point benchmarks, it is 13% slower. (All results except for Norton SI were obtained using systems with a 256-Kbyte, write-back cache. For Norton SI, the DX-33 and DX-50 systems used a 128-Kbyte, write-through cache.)

Comparisons with the DX-25 are more difficult, since Intel does not include results for this clock frequency in its current benchmark report, and earlier benchmark reports used different compilers, benchmark versions, and system configurations. Intel characterizes the DX2-50 as being typically 70% faster than a DX-25, although this is highly dependent on the system design and the application.

Because the DX2 has a much higher bus utilization than the standard 486DX, it is more sensitive to the performance of the external cache and memory system. Cacheless DX-25 systems with a good memory system design perform nearly as well as cached DX-25 systems. When upgraded with a DX2-50, however, the performance difference between the two systems will increase. Intel's benchmark report shows that adding a 128-Kbyte write-through cache to a DX2-50 system increases performance by an average of 10% for DOS ap-

plications and 16% for Windows applications. Even with the cache, reducing the DRAM write latency by one clock cycle boosted Dhrystone performance by 24% and SPECint89 performance by 13%, illustrating the importance of an optimized memory system.

The Upgrade Market

The same chip will be offered later this spring in a slightly different pinout as an end-user upgrade product. This version, code-named P23T and to be called the OverDrive processor, is pin-compatible with the 487SX—the 486DX in a modified pinout sold as an upgrade for 486SX systems. Intel has been encouraging makers of 486SX systems to include the OverDrive

socket, formerly called the Performance Enhancement socket, in preparation for the launch of the clock-doubler processors. There is no reason why 486DX systems cannot include the same socket, and now that the OverDrive processor is about to become available, some vendors will add this socket to their systems.

The OverDrive pinout differs from the standard 486DX in that it includes an alignment pin to make it harder to insert the chip incorrectly, and one signal (FERR#) has been arbitrarily moved to a different pin. Intel apparently wants the various 486 versions not to be pin-compatible so it more easily can pursue different pricing and marketing strategies for the different versions (prompting one of our contributors to call this scheme "pins for profit").

When an OverDrive processor is installed, the original processor is electrically disabled. Intel discourages users from physically removing the original processor, in part because of potential damage to the system board and in part because they don't want to create a supply of used 486 chips. Once an OverDrive processor is installed, however, the original processor can be removed, and savvy dealers and users are likely to do so if they can find a market for the used chips.

The OverDrive processor is the key product that Intel's massive "vacancy" advertising campaign is intended to promote. This ad campaign may well be Intel's least-effective use of its cash since its investment in BiiN. Some of the advertisements don't even mention the clock-doubler concept, and they are intended to promote the general concept of plug-in processor upgrades. Many users have assumed that these advertisements were for the 487SX, since that has been the only announced product for this socket. As a result, these users view these ads as a desperate act on Intel's part to attempt to sell them a chip that makes little sense. Now that the OverDrive chips are about to be announced, the

Price & Availability

The 486DX2-50 is in production now and is priced at \$550 in quantities of 1000. The 486DX2-66 will be available in the second half of the year; pricing has not been released.

For more information, call 800/548-4725 or write to Intel, P.O. Box 7641, Mt. Prospect, IL 60656-7641 and ask for literature packet #JB-47. The order number for the 486DX2 Performance Brief is 241254-001.

ads will presumably become more specific, but so far they have been of questionable value.

The 486DX2 can be used to upgrade some 486DX systems by removing the CPU and replacing it with the DX2, but there are several potential problems:

- The chip's power consumption is substantially higher, so the cooling in some systems may not be adequate. The DX-25 draws 550 mA typical and 700 mA maximum, while the DX2-50 draws 775 mA typical and 950 mA maximum.
- While the interface timing *specifications* are identical, the *actual* timing is slightly different. This can cause problems in some marginal system designs.
- Some BIOS programs include speed-dependent timing loops.

Intel says that its testing revealed that about one system in four encountered problems. Making a list of systems that can be safely upgraded isn't as easy as it might seem, since it sometimes depends on which revision of the system board and BIOS is present. Some system vendors may certify their systems as upgradeable. Computer dealers may offer unauthorized upgrades, and sophisticated end-users may be willing to try the upgrade themselves, but the potential for problems is significant.

From the system maker's perspective, upgrade chips cut them out of the upgrade sale. Thus, some may prefer to sell upgrade CPU cards, offering larger caches or other features in addition to the faster processor. Some system makers more focused on their profits than on benefits to their customers might *ensure* that their BIOS contains speed-sensitive code as a way of making users come to them for an authorized upgrade, so they could charge more than the street price of a 486DX2 chip for the new processor and a new BIOS ROM.

Conclusions

Intel's 486DX2-50 will give system makers an easy way to upgrade their 25-MHz system designs, providing a cost-effective alternative to true 50-MHz systems. When the 486DX2-66 becomes available, it will provide a new pinnacle of 486 performance without stressing the ability of PC makers to build systems around it. The end-user versions of these processors—the OverDrive chips—will expand Intel's market by giving them a chance to sell a second processor to 486 system buyers.

This technology will also be important for Intel because it will help keep them one step ahead of 486 cloners. With the 386, Intel let its design age without upgrading it, making it easy for other semiconductor vendors to provide faster versions. With the 486, Intel will not give them this opportunity, and the clock-doubler chips are one more weapon in this battle. ♦

Rambus

Continued from page 18

RDRAMs also threaten to replace VRAMs in frame buffers. With the higher bandwidth of the RDRAMs, it is practical to use part of main memory for the frame buffer. This would have some impact on the CPU performance, since transfers from the frame buffer to the video output logic would contend with CPU requests. For maximum performance, the frame buffer controller could have its own private Rambus with a dedicated frame buffer. The separate video port provided by VRAMs would not be needed, since the Rambus interface provides more than enough bandwidth for processor access and screen refresh, even in a high-resolution, full-color system. Unlike VRAMs, RDRAMs are not specific to video applications, so they have the potential for higher volume and the resulting lower costs. RDRAM chips will also be smaller than VRAMs.

Conclusions

Rambus has high expectations for its technology—the company predicts that RDRAMs will account for over 50% of DRAMs within five years. While it is conceivable that this degree of success will be realized, there are several obstacles.

The price of RDRAMs is a key issue. Rambus claims that RDRAMs are about 20–25% larger than conventional DRAMs for 4-Mbit devices, with the area premium dropping to about 10–15% for 16-Mbit chips and 5% for 64-Mbit chips. Rambus DRAMs will initially have a much lower volume than conventional DRAMs, however, and there will not be as much competition, so the price premium could be higher.

Some designers are skeptical of whether Rambus will meet its 250-MHz clock rate goal, but Rambus says that test chips have been fabricated and they have had no difficulty reaching this rate. Until RDRAMs and systems using them are in production, however, some question will remain regarding the practicality of this speed. Another question is whether the high hit rates will be achieved in systems with only a few RDRAMs; as the hit rate drops, performance drops drastically.

There may also be some resistance among system vendors to using a technology that requires royalties to be paid to Rambus. This royalty will be buried in the cost of RDRAMs, interface ASICs, and eventually processors, but it is nonetheless present. System makers will also have a limited choice of vendors for interface ASICs, since they must go to a vendor that has licensed the Rambus interface cell.

If these barriers are overcome, however, the Rambus technology has the potential to reshape memory system design, graphics system architecture, and microprocessor bus interfaces. ♦