MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 7 NUMBER 6 MAY 10, 1993

AMD Jumps Into 486 Market

Buoyed by Legal Upset, AMD Begins Shipping 486 with Intel Microcode

By Michael Slater

Encouraged by a ruling that threw out the jury's decision in the 287 microcode case, AMD has begun shipping 486 microprocessors with Intel's microcode. AMD has been ready to ship the chips since last fall, but was stymied by its loss in the key microcode case. The company has been building the chips since then for internal testing purposes and as a vehicle for fine-tuning the process. As a result, it is now able to move the chip into production rapidly and even has some stock (which the company declined to quantify, other than as "small") from which to begin shipments.

The overturning of the verdict had not been anticipated because it is extremely rare for a judge to throw out a jury verdict due to withheld evidence. AMD expects that its chances of winning on the next trial are significantly enhanced by the fact that it has the new evidence and has seen Intel's defense. Furthermore, many observers believe that Intel's victory in the first trial was more a result of confusion among the jurors than any reflection of the facts.

AMD said that it already has customers for the chips, but the company declined to name them. While AMD's 486 does not have the clock-rate advantage over Intel's chip that its 386 enjoys, AMD has established dozens of customer relationships with PC makers and a track record for providing compatible products that should make 486 sales easier than its early 386 sales.

Intel has reportedly not been able to meet all the demand for the 486, so there should be a ready market for additional supply, especially considering the relatively limited volumes AMD can provide in the near term. AMD has said that it will ship "thousands" of chips in May, with production ramping throughout the rest of the year. Production is expected to reach 1 million units per quarter in 1994. AMD projects that it will achieve a 5% market share by the end of 1993, and a 20% share by the end of 1994 (both based on run rate at the end of the year). The latter figure is dependent on an anticipated foundry agreement to increase AMD's capacity.

Whether AMD has a legal right to use Intel's microcode has still not been established; Judge Ingram's ruling simply determined that the issue would be retried. By introducing its 486 with Intel microcode, AMD is gambling that the outcome of the next trial will be different from the first. If it loses this gamble, the company would be liable to Intel for whatever profits it earns on its Intel-microcode 486 (or, at Intel's discretion, whatever profits Intel can show it lost as a result of AMD's sales).

In a recent interview with *Microprocessor Report*, AMD CEO Jerry Sanders conceded that the company made a mistake in basing its initial plans on a positive outcome from the microcode litigation and not investing in clean-room microcode from the start. He explained his confidence by saying, "I was there. I know what we negotiated. The intent [of the disputed agreement] was to allow us to produce Intel-compatible microprocessors."

AMD still plans to introduce its clean-room microcode version in July (symbolically, on the fourth of July), and production eventually will be phased over to that version. Customers will want to separately qualify the clean-room part, but they can continue using Intel-microcode chips while this qualification is in process. If the clean-room microcode is fully compatible and is completed on time, then the effect of the Intel microcode part will be just a three-month acceleration in shipments.

Should there be any delays or problems with the clean-room microcode, however, AMD's ability to use the Intel microcode could be significant. It is certainly easier to convince prospective customers of the compatibility of the part with Intel microcode, which is a key reason why AMD pursued this path in the first place. AMD may find it tough to get customers to switch to the AMD microcode part if they have the option to stay with the Intel microcode version, which would increase AMD's exposure to a future legal loss.

AMD Starts with DX and DX2

As with its 386, AMD's 486 is derived from Intel's logic design but has been modified to provide static operation. Fabricated in a 0.7-micron, three-level-metal pro-

cess (called CS14), the chip is 89 mm^2 (138K mils^2)—slightly larger than Intel's implementation, which is 84 mm^2. Since the chips are derived from Intel's logic design and Intel's microcode, they should be state-for-state compatible and provide identical performance at a given clock rate.

AMD initially will offer only the 486DX and 486DX2, deferring the lower-profit 486SX until it has more production capacity (see sidebar). AMD says it is shipping 33- and 40-MHz versions of the 486DX "from stock." Power consumption, at 600 mA typical (700 mA maximum) for the 33-MHz part, is 100 mA less than Intel's because Intel still builds its 33-MHz chip on its 1-micron process. Power consumption at 40 MHz is 700 mA typical, 850 mA maximum.

The 486DX2-50 is now sampling and is promised for production later this quarter, and the 486DX2-66 is planned for production in the third quarter. None of these chips offer any advantages over Intel's; they are simply plug-compatible replacements (except that Intel does not offer a 40-MHz chip). Power consumption is the same as for Intel's chip, at 775 mA typical.

AMD is exactly matching Intel's prices: $306 for the 486DX-33, and $417 for the 486DX2-50. Following its 386 strategy, the 486DX-40 is being offered at the same price as the 486DX-33. As long as AMD is production-limited, it will seek to keep the price umbrella up. While 486 prices are likely to drop significantly in the long run as a result of AMD's introduction, the big drops will probably not occur until 1994 when supply begins to exceed demand.

The 40-MHz 486DX is a speed that Intel skipped in an effort to limit the proliferation of 486 speed versions. Intel's 50-MHz chip was the first 486 built in its 0.8-micron process. The new process apparently provided adequate yield at 50 MHz, so there was little incentive to add a 40-MHz version. The 50-MHz chip has fallen into disfavor in the market, however, because the 486DX2-66 offers higher performance for most applications and an easier system design, since it has a 33-MHz bus.

AMD's 40-MHz part offers makers of 486DX-33 systems an upgrade alternative to the 486DX2-66. AMD will not charge any premium for the 40-MHz part, while the DX2 chips are significantly more expensive, so this will be a less-costly enhancement. Several chip-set vendors have already announced plans to support the 40-MHz chip (see Price & Availability sidebar). AMD expects that the 486DX-40 will gradually obsolete the 486DX-33, just as its 386DX-40 all but eliminated the 386DX-33 from the marketplace.

One key difference, however, is that Intel isn't likely to allow itself to be left out this time. Because Intel was focused on the 486, it did not offer the 40-MHz speed grade of the 386. (This was presumably a marketing decision, but Intel was also hampered by the fact that it did not move the 386 to its 0.8-micron process.) With the 486, Intel has been shipping a 50-MHz part for some time, and if AMD succeeds in moving much of the 33-MHz market to 40 MHz, Intel could add this version to its line-up with nothing more than a change in the chip label.

Little Differentiation at First

AMD's first differentiated version is the 486DXLV-33, a 3.3V version that offers static operation, a system-management mode for power management, and a 196-pin PQFP package. This chip is sampling now, with production planned for July, at the same pricing as the standard 33-MHz part.

AMD's 486DXLV is very similar, in concept, to Intel's rumored S-series—enhanced 486 chips that are expected to offer the same list of features as in AMD's 486DXLV and should ship soon afterward. The primary difference, based on what is currently known about the as-yet-unannounced S-series chips, is that AMD's and Intel's SMM modes differ in the details of their operation. It remains to be seen whether either processor will

Capacity Issues

Aside from the legal challenges, capacity is a key issue for AMD; the only facility at which it can currently build the 486 chip is its sub-micron development center (SDC) in Sunnyvale. AMD began a $160 million campaign at the end of last year to outfit this facility as a production fab for 486 processors and flash memory. When fully outfitted in mid-94, it will have a capacity of nearly 3,000 six-inch wafers per week, and AMD has stated that it expects to ship $250 million worth of 486 chips in the first 12 months of production and to achieve a run rate of $100 million per quarter from the SDC alone. AMD is seeking outside foundry capability—reportedly including IBM and Hewlett-Packard—to boost its 486 capacity.

The SDC is currently being used for flash memory production, but AMD is in the process of moving flash production to its Fab 14 in Austin, Texas. The SDC is also used for research and development and for some 29000-family production. Based on our estimate of at least 60 good die per wafer and a low estimate of $150 average selling price, AMD would need less than 1,000 wafers per week to reach its $100 million quarterly goal.

AMD's ability to ramp its 486 capacity beyond this level—and to build chips using its 0.5- and 0.35-micron processes currently in development—is dependent on an as-yet unbuilt plant, called Fab 25. Groundbreaking for Fab 25 is scheduled for this June, with early production expected by the end of 1994 and full production in 1995. This facility, adjacent to AMD's current buildings in Austin, initially will include 60,000 square feet of clean-room space. It will be capable of producing 5,000 eight-inch wafers per week when fully built-out to its 80,000 square foot capacity.

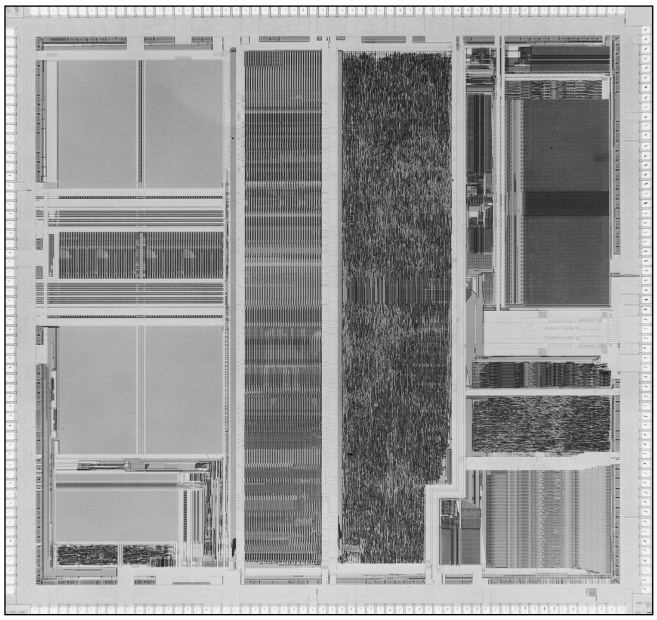

Die photo of AMD's 486, which measures 91 × 97 mm (360 × 384 mils) and includes 930,000 transistors. (Intel claims a transistor count of 1.2 million for its 486, but this counts unused transistor sites in the microcode ROM and PLAs, which AMD does not count.)

have a significant advantage in power consumption.

AMD is likely to seek ways to further differentiate its parts from Intel's, and the clearest opportunity is in the caches. Changing the caches to a write-back instead of a write-through design would provide a noticeable performance improvement—certainly enough to put systems using such a chip at the top of the charts in magazine roundups, which is a key factor in PC sales—and would have minimal additional cost.

Such a cache cries out for burst writes (for dirty cache line write-backs), which the 486 bus does not implement, so AMD might enhance the bus in this way. Bus extensions would also be needed to support cache coherency; a write-back cache must be snooped on read cycles from other bus masters, while a write-through cache needs to be snooped only on write cycles. (Chip-set makers are already revising their designs to support write-back caches for Intel's Pentium and the future 32-bit-bus version, the P24T.)

The additional signals could be added on "no-connect" pins of the standard 486, and the chip could default to write-through mode, providing full compatibility with existing designs. To fully exploit a write-back cache, however, some system design changes would be needed, so it is natural for AMD to delay introduction of such a part until it has established its presence in the 486 marketplace and tapped into the easiest business—simply filling unmet demand for standard 486 chips.

Doubling the cache size is another opportunity. This would be especially useful for the 486DX2 chips, since the higher internal clock rate doubles the cache miss penalty. Chips with larger on-chip caches could be fully pin-compatible with standard 486 chips, and as with the write-back cache, the larger cache would enable systems using the chip to lead any performance comparisons. Such a device would require a new chip layout, of course, and its larger size would increase manufacturing cost, so AMD isn't likely to introduce such a chip immediately.

If Intel should offer a 486 with a larger cache, however, it would put considerable pressure on AMD to do so. With Intel moving its 486 to a 0.6-micron process next year, a larger cache is likely—especially since Intel is aware that this would otherwise be a differentiation opportunity for AMD.

There are numerous other possibilities for products based on the 486 core, including clock-tripled versions (such as a 486DX3-99), chips without floating-point units (for example, Intel doesn't offer a 486SX2), chips that plug into Intel's OverDrive socket, and products like Cyrix's 486SLC/DLC that combine a 486 core with a 386 bus interface. AMD declined to comment on what versions it might offer in the future, but the company clearly has plans to broaden its product line and could, in theory, produce whatever of these products make marketing and financial sense.

Price & Availability

The Am486DX-33 and Am486DX-40 are in production now and are priced at $306 in thousands. The Am486DX2-50 is sampling now, with production in June, priced at $417 in thousands. The Am486DXLV-33 is sampling now, with production planned for July, and is priced identically to the Am486DX-33.

OPTi, PicoPower, Western Digital, Symphony, and ETEQ all plan to offer chip-set support for the 486DX-40. The first three companies also plan to support the 486DXLV's system management mode for power management. Phoenix Technologies will provide BIOS support for power management.

Advanced Micro Devices, PO Box 3453, Sunnyvale, CA 94088; 800/222-9323 or 408/749-2036.

Legal Issues Still Unresolved

AMD's sudden entry was precipitated by a favorable ruling from Judge Ingram, issued in response to a motion from AMD alleging that Intel withheld evidence from last year's jury trial regarding AMD's right to use Intel's microcode in its 287 math coprocessor. (Although the case was focused on the 287, it affects the 486 because it establishes how the Intel/AMD agreement should be interpreted.) The judge agreed with AMD's accusations and therefore granted AMD's motion for a new trial. No trial date has been set, and none is likely until the fall at the earliest.

The dispute centers on the interpretation of a 1976

386 Still Thriving—For Now

Despite all the press about the demise of the 386 market, AMD continues to ship them in record numbers. Much of the system demand is coming from countries other than the US, including Eastern Europe and other developing areas where the 386 vs. 486 price difference is still significant.

In the first quarter of 1993, AMD shipped 3.15 million 386 chips—more than in any previous quarter. Revenues from 386 sales dropped over the course of 1992 because prices dropped; unit volumes have been steadily increasing *(see **070101.PDF**)*.

AMD expects to ship in excess of 3 million 386 chips in the second quarter, but it is not yet clear whether it will exceed the first quarter shipments. Volumes of 386 chips are expected to begin dropping in the third quarter. While AMD's 1993 unit shipments of 386 chips might not be far below 1992 shipments (which totalled over 9.5 million units), revenue from these shipments will be significantly lower; 1994 volume is likely to be relatively low, and at even lower prices, making the 486 essential to AMD's financial picture.

The 386DX represents two-thirds of AMD's current 386 shipments, with essentially all of the 386DX volume at the 40-MHz clock rate. Over half of AMD's 386SX shipments are at 33 or 40 MHz. High-volume pricing is below $30 for the SX and well under $50 for the DX. AMD expects pricing for the SX to drop below $20 in 1994.

agreement between Intel and AMD that was intended to give AMD the right to manufacture Intel microprocessors as an alternate source. The agreement clearly gives AMD the right to use Intel's patents, and this has not been disputed. The subject of the dispute is whether the agreement also gives AMD the right to Intel's copyrighted microcode, which is covered in a separate section of the agreement. *(see **060901.PDF** and **0616MSB.PDF**)*

Before the case went to trial, Intel publicly argued that the agreement gave AMD the right to copy the microcode, but not to distribute it, because that is what the agreement literally said. At the trial, however, Intel dropped this line of reasoning and instead argued that the term "microcode" referred to monitor programs and other software in external ROM, not on-chip microcode. The reasoning was that the agreement was intended to allow AMD to build development systems for its own use; thus, this was called the "blue box" argument. In support of this view, Intel pointed out that none of its microprocessors had microcode at the time the agreement was signed. Nevertheless, the two companies were no doubt aware that forthcoming parts would have microcode, and it is AMD's view that the intent of the agreement was to allow AMD to produce these parts.

There are four documents that Intel did not provide to the court that were the basis for overturning the jury verdict: an August, 1990 "Litigation Reporter," which is an internal document provided to senior management and public relations employees to assist in explaining Intel's position to the public; two math coprocessor competitive summaries; and an Intel news release discussing the litigation.

These documents all reflect Intel's initial stance that the agreement gave AMD the right to copy, but not distribute, the microcode. For example, quoting from the Litigation Reporter discussing the 287 situation:

> "Intel believes that the license does not permit AMD to copy and distribute Intel's microcode. (It allows them to 'copy' or use the microcode, but not to "DISTRIBUTE"). In doing so they infringe Intel's Copyrights."

The parenthetical statement quoted above directly contradicts the core of Intel's argument at the trial that the agreement did not refer to the microcode in the 287.

Intel claims that its two different explanations for why AMD couldn't use the microcode are not contradictory, but are merely "two prongs" of its attack. Intel says that its public statements focused on the "copy, but not sell" approach because it was simpler to explain to lay people. AMD, on the other hand, calls Intel's claim that the term "microcode" in the agreement did not refer to the microprocessor's microcode a "fabrication" and a "made-for-trial" argument.

Intel argued that AMD had access to all non-disclosed information from public sources, but AMD countered that such information did not have the same "evidentiary impact" as a corporate admission from Intel.

Intel's most egregious action, according to AMD attorney Rich Lovgren, was its alteration and misrepresentation of a November 1991 Litigation Reporter. After the document was sent from Intel to its attorney's office, the cover memo and a distribution list were removed before it was submitted to the court. These were the only dated pages. In court testimony, Intel's witness mistakenly identified the document as being from 1990 instead of 1991. As a result, AMD was unaware that Intel had failed to produce the document it was seeking—the November 1990 Litigation Reporter.

In the end, the judge concluded that,

> "[the four documents] would all be admissible as party admissions and should have been produced before trial. The failure to produce them substantially interfered with AMD's discovery and trial presentation. Moreover, the non-production prevented AMD from fairly presenting its defense."

More Legal Challenges to Come

The legal bickering is far from over, of course. Intel

has asked the judge to reconsider his overturning of the jury verdict. Since there is little chance that the judge will do so, Intel has asked the judge to "certify" his ruling so that Intel can appeal it without waiting for a second trial to be held.

Intel has also filed a copyright infringement suit against AMD's 486, challenging the legality of both the Intel-microcode and the clean-room versions. In addition to repeating the copyright issue as tried in the 287 case, Intel repeated a claim it made regarding the 386: that the PLA contents are software (and thus protected by copyright) but not microcode, so they are not covered by the agreement. Intel alleges that both the "Overall Control Program" and the "Floating-Point Control Program" fall into this category. These claims break new legal ground; it is unclear whether PLA contents will be accepted as software, since they can be viewed as a representation of a logic design, which cannot be copyrighted.

Intel also claims that AMD has used circuitry and microcode designed for support of Intel's in-circuit emulators to implement its system management mode. This is an issue because the Intel/AMD agreement explicitly prohibits AMD from producing "bond-out" versions of the parts that provide access to this circuitry and microcode.

Intel has also challenged AMD's right to have the chips made by a foundry, asserting that AMD's alleged right to copy does not extend to a right to "have copied."

Even though AMD has not yet shipped its clean-room microcode version of the chip, Intel has already challenged its legality. Intel claims that AMD's copying of the Intel microcode for disassembly constitutes copyright infringement and therefore taints the resulting clean-room code. This is the same argument Sega used unsuccessfully against Accolade *(see **061605.PDF**)*, but Intel argues that the Sega ruling allows disassembly only if it is the only way to determine the required functions. According to Intel, the functions of the 486 microcode can be determined by studying the data sheet and observing the operation of the chip, so the conditions of the Sega case do not apply.

Limited Impact in '93

AMD will now be in the 486 market a few months earlier than expected, and with a chip whose compatibility will be much harder to challenge. AMD has taken a risk that it might still lose in court, but the company is confident of success, and it plans, in any case, to switch to a clean-room design—though its customers may not be as enthusiastic about the change. Intel asserts that whatever copyright license AMD holds (but not the patent license) expires at the end of 1995; AMD disagrees with this interpretation, and this issue has not yet been heard in court. If Intel were to prevail on this point, AMD would be forced to switch to clean-room microcode by 1996.

The effect on Intel is likely to be small this year, since the market demand has outstripped Intel's ability to supply the 486 and AMD's production capacity is limited. AMD will not be shipping the fastest—and most profitable—486 chip, the 486DX2-66, until later in the year. AMD is therefore unlikely to cause Intel's shipments to decrease in the near term, and any significant impact on pricing is many months away. In the long run, however, the 486 business will become much more competitive, supply will exceed demand, and prices will inexorably plummet. The acceleration of AMD's entry into the 486 market might encourage Intel to speed up its Pentium production ramp, but the vastly higher production cost of Pentium will prevent Intel from using it as a "486 killer" any time soon.

One way in which the court decision, if it is ultimately decided in favor of AMD, could be significant is with regard to Pentium. A decision in favor of AMD would enable AMD to produce an Intel-clone Pentium chip, complete with Intel microcode, if it should choose to do so. Currently, AMD's plan is to create its own independent design for this performance level, and this design is well under way. AMD expects to have a Pentium-performance processor in production in 1995.

As Sanders gleefully pointed out, the end of Intel's 486 monopoly is good for all companies in the industry—except one. For the remainder of 1993, the actual effects on Intel are likely to be minor, although the press and the stock market may exaggerate them. In 1994, however, the 486 marketplace will have an entirely different character. ♦

