

Big Plans for Small WinChip

IDT to Add 3DNow, Integrated Chip Set, Superpipelining

by Linley Gwennap



Although IDT has shipped only a few hundred thousand x86 processors, the company has big plans for the future. At the recent PC Tech Forum, Glenn Henry laid out plans for several future versions of the WinChip processor. Henry, founder and head honcho of IDT's Centaur design subsidiary, described chips with 3D extensions, with larger L1 caches, with an integrated north bridge, and with a high-speed superpipelined design, as Figure 1 shows.

All of these processors are aimed at the low end of the PC market—systems selling for less than \$1,000. This relatively small but fast-growing segment should provide plenty of opportunity for IDT's wares. The WinChip should have a manufacturing-cost advantage in this segment, allowing IDT to underprice its competitors. We expect the company to ship more than one million x86 chips in 1998, and up to 2–3 million the following year.

WinChip 2 3D Improves Original Design

The company is currently shipping its original WinChip (also known as the C6) at clock speeds of 200, 225, and 240 MHz. The last version is good only for those who value clock speed above all else; because it uses a 60-MHz bus, the 240-MHz WinChip actually delivers less performance on many PC applications than the 225/75-MHz version. The original WinChip uses a relatively simple scalar design (see MPR 6/2/97, p. 1) to deliver integer performance similar to that of a Pentium/MMX at the same clock speed.

That chip, however, is much slower than Pentium on applications that make heavy use of FP or MMX instructions. To get its first x86 chip to market in record time, the Centaur team focused on integer performance and x86 compatibility. For the second version, called the WinChip 2 3D, the team went back and added some features that close the performance gap.

According to IDT, the new design delivers MMX performance similar to Pentium/MMX's. Like the Intel chip, the

WinChip 2 3D can issue two MMX instructions per cycle, and it has similar latencies. The original WinChip was hampered because it could issue only one MMX instruction per cycle. The FPU was completely redesigned to be fully pipelined with latencies similar to Pentium/MMX's. According to IDT's measurements, the WinChip 2 3D should match Pentium/MMX on most FP applications.

The design contains other enhancements (see MPR 11/17/97, p. 17) that improve performance on integer applications. In addition, IDT adopted the 100-MHz version of Socket 7 being promoted by AMD (see MPR 6/1/98, p. 16). Compared with the original WinChip, the improved core provides a 10% boost on Winstone 98, an integer-only benchmark. This boost is a bit larger than the company had anticipated.

The biggest change from the previous disclosure is support for 3DNow (see MPR 6/1/98, p. 18). Henry had originally planned to include a set of proprietary 3D extensions that went well beyond 3DNow by adding 22 new registers and 53 new instructions. After AMD's announcement, however, he realized that it would be better for the companies to

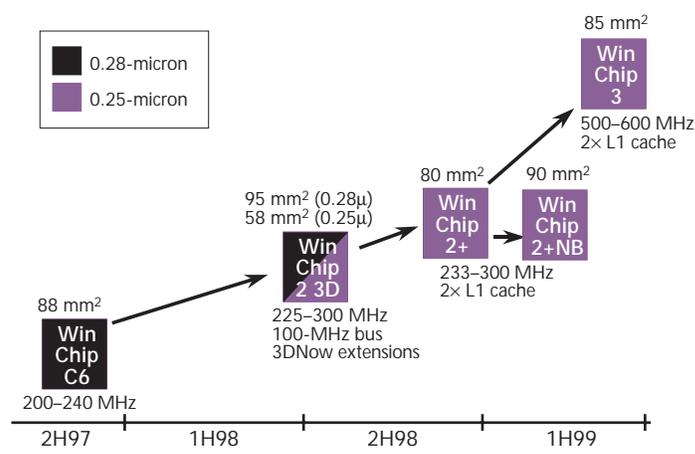


Figure 1. IDT's WinChip roadmap includes processors with 3D extensions, with an integrated north bridge, and with a high-speed superpipelined design. (Source: IDT)

join forces behind a common standard than to promote incompatible extensions. Changing WinChip's 3D extensions caused significant redesign, but the amazingly fast Centaur team kept the schedule delay to just a couple of months.

The WinChip 2 3D is now sampling; IDT expects volume shipments in July. The initial 0.28-micron version measures 95 mm², just 7 mm² larger than the original WinChip. Henry said that most of the extra area came from adding dual-issue logic for the MMX unit. The new 3D unit added only 2 mm², and the faster FPU fits in the same area as the previous FPU. Figure 2 shows the new layout.

Centaur cut one corner to keep the 3D unit small. On most instructions, the WinChip 2 3D is as fast as AMD's K6-2 (see MPR 6/1/98, p. 16). The reciprocal square-root (PFRSQRT) instruction, however, takes 10 clocks to complete, versus 2 clocks for the AMD chip. Henry said matching AMD's performance on that function would have greatly increased the size of the 3D logic. Reciprocal square root is typically used in 3D lighting calculations, so the WinChip 2 3D is likely to trail the K6-2 in performance on 3D games.

Enhancements to the WinChip 2

IDT already has prototypes of the WinChip 2 3D in its 0.25-micron process. By the end of the year, the company expects to have switched its WinChip production entirely to the new process. At that point, IDT's foundry arrangement with IBM should also kick in. IBM offers additional 0.25-micron capacity and, perhaps more important, an Intel patent license.

The 0.25-micron process offers only a minor speed boost, from 266 to 300 MHz, due to the small decrease in transistor size. The metal layers get a much greater shrink,

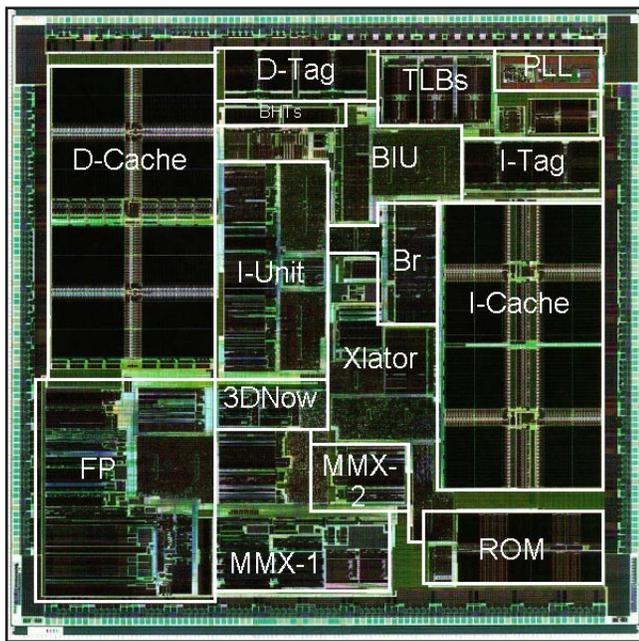


Figure 2. IDT's WinChip 2 3D measures 9.7 × 9.7 mm in a 0.28-micron four-layer-metal process. The chip is sampling; this die plot shows the lower layers, which are not visible in a die photo.

which will reduce the WinChip's die size to a tiny 58 mm². This is 30% smaller than the K6-2 and less than half the size of the Deschutes CPU used in the Celeron processor. It is smaller than even the 0.25-micron Pentium/MMX, which Intel doesn't offer for desktop PCs. According to the MDR Cost Model, the 0.25-micron WinChip will cost about \$28 to build, less than the cost of any competitor's part.

The current WinChip (and the first WinChip 2 parts) supports only integer clock multipliers. This prevents the chip from running at 233 MHz with a 66-MHz bus, for instance. The 0.25-micron version will add a new clock circuit that supports bus multiples in halves and thirds. Using multiples in thirds is unusual, but it has the advantage of allowing the CPU to run at familiar speeds such as 233 and 266 MHz while keeping the system bus at its full 100-MHz speed.

IDT is also quietly working on a mobile version of the WinChip 2. The chip supports a split supply, with the core running at 2.5 V and the I/O at 3.3 V. In this mode, the 0.25-micron WinChip will dissipate about 6 W, well within the notebook power range. Notebook vendors don't like the standard PGA package used by Socket 7 processors; Henry said the company is working on a "special" package for notebooks but wouldn't comment on when it might be available. We expect the notebook version to use a BGA package.

Integration Options Lead to North Bridge

Last fall, IDT announced plans to integrate a 256K level-two (L2) cache onto the WinChip in 2H98. This plan seemed natural for IDT, given its extensive SRAM experience. After further analysis, however, the company has changed its plan. According to Henry, integrating the L2 cache would have added 55 mm² in the 0.25-micron process, nearly doubling the size of the chip! Depending on the application, the performance gain would be 0–20% over an external 256K cache, yet these SRAMs add only a few dollars to the system cost.

Instead, IDT will simply double the size of the on-chip level-one (L1) caches, which will grow to 64K each for instructions and data in the WinChip 2+. Simulations show this change will deliver a similar performance boost while adding only 20 mm² to the die. In taking this approach, IDT is flying in the face of forthcoming chips such as Intel's Mendocino and AMD's K6-3, both of which integrate a sizable L2 cache without enlarging the L1 caches.

One reason some vendors may be shying away from larger L1 caches is timing. At high clock speeds, accessing a large cache in a single cycle can create a critical timing path. At 300 MHz, however, the WinChip 2 isn't being clocked as fast as other 0.25-micron processors, giving it a bit more leeway. Henry says the L1-cache access is not a critical timing path on the WinChip 2.

Instead of devoting extra die area to the L2 cache, IDT decided to push further and integrate the north bridge of the chip set instead. In addition to the usual advantages of integration (lower cost, smaller footprint, lower power), this design offers a performance advantage. Connecting the pro-

cessor directly to main memory allows the CPU to access the DRAM more quickly. This direct access eliminates bus synchronization and other overhead, improving application performance by reducing the number of cycles that the CPU has to wait for main memory.

Other vendors are pursuing different integration strategies. Intel's forthcoming Whitney chip set (see MPR 4/20/98, p. 18) combines the north bridge and graphics chip. Cyrix's MediaGX and forthcoming MXi combine the CPU, north bridge, and graphics chip. Both of these alternatives restrict OEM flexibility by integrating the graphics function. IDT's combo chip, dubbed the WinChip 2+NB, still permits the use of any AGP-based graphics chip.

The downside of integration is that the WinChip 2+NB is no longer a Socket 7 part; it requires a new motherboard design. The redesign should be fairly easy, as the integrated chip has essentially the same pinout as a standard north bridge; the CPU is simply hidden inside. The WinChip 2+NB has the same interfaces as a Socket 7 north bridge, including a Pentium CPU bus that is used to connect to the external L2 cache.

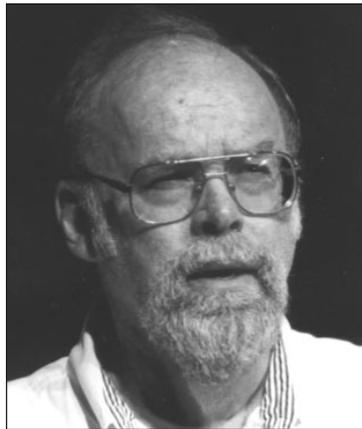
IDT is working with a system-logic vendor to acquire the north-bridge design. The company wouldn't specify its partner, but presumably it is either VIA, SiS, or Acer Labs, which have all announced Socket 7 chip sets with AGP. The WinChip 2+NB chip is scheduled for volume shipments in 1Q99, one quarter after the Socket 7 version of the WinChip 2+.

Superpipelining Maximizes MHz

Henry continues his contrarian approach when looking to the next CPU generation. All other x86 vendors have adopted superscalar designs to compete with Intel's P6. The P6 and AMD's K6 go further by reordering instructions to execute several per cycle.

Henry believes this is folly. Given the relatively small performance gains delivered by these complicated devices, he may have a point. In the WinChip 3, Centaur plans to maintain its scalar pipeline while dividing each pipe stage in half. Henry expects this new design will double the CPU's clock speed, allowing it to reach 600 MHz in a 0.25-micron process. This projection doesn't seem to take into account delays from latch overhead, which may force IDT to move to a more advanced process to reach 600 MHz.

This relatively simple design technique does not double performance, because pipeline penalties and memory delays are also doubled. Henry expects to see a 30% improvement on Winstone 98 and up to 80% better performance on certain CPU-bound benchmarks. Yet the die-size impact of superpipelining the chip is only 5-10%, far less than adding superscalar execution or instruction reordering.



Centaur president Glenn Henry outlines the next generation of WinChips at PC Tech Forum.

THE PHOTO GROUP

Price & Availability

IDT is now sampling the WinChip 2 3D, with volume production expected in July. The company has not yet announced pricing for the new part. For more information, access www.winchip.com.

A key benefit of superpipelining is that the WinChip 3 could become the fastest 0.25-micron x86 processor available, at least in terms of clock speed. Many buyers, particularly in the low-end space that IDT targets, pay little attention to benchmarks and focus mainly on MHz. The WinChip 3 will be an ideal product for these buyers.

IDT has not yet decided on the bus interface for the WinChip 3. It will probably stick with Socket 7, but the company is also investigating both Intel's Slot 1 and AMD's future Slot A. IDT plans to offer a WinChip 3+NB as well, getting around the bus issue completely. If the integrated north-bridge concept catches on, the company may not bother with the standalone WinChip 3.

Given the skill of the Centaur team and the relatively modest changes required to superpipeline the chip, IDT hopes to ship the WinChip 3 by mid-1999. In the 0.25-micron process, the chip should measure about 85 mm² and reach 600 MHz, according to Henry. About the same time, the company expects to have a 0.22-micron process available. This process could boost clock speeds to 700 MHz or higher while pushing the die size below 60 mm² again.

Aiming High at the Low End

For a small player in the x86 market, IDT will be very busy, with five new designs scheduled to be released in the next 12 months. Most of these designs are fairly minor variations, so the plan appears aggressive but doable. If all goes well, WinChip 3 should deliver performance similar to Intel's midrange processors. Even if IDT doesn't meet this goal, its processors should have plenty of performance for the low-end markets it is targeting.

IDT's biggest challenge now is to demonstrate it can deliver parts in volume. Then it can begin signing larger customers and work its way up to the top tiers. The WinChip's biggest advantage is price: IDT can afford to sell it for less than any competitor's product. For this reason alone, the company should be able to sell as many chips as it can make for at least the next year. We expect IDT to make life difficult for AMD and Cyrix at the lowest end of the market. 