x86 Architecture Fragmenting

Each Chip Vendor Creating Different Extensions to Speed 3D Geometry

After years of stability, the x86 instruction-set architecture has entered a period of relatively rapid change. Surprisingly, Intel's competitors are out in front, making extensions without waiting for Intel to set the standard. The resulting proliferation of extensions could make it difficult for the chip vendors to get the software support needed to deliver value to users from these innovations. The bigger question these extensions raise is whether any company but Intel can extend the architecture in a meaningful way.

Intel has made some efforts to keep the architecture unified. Intel did not try to use MMX as a competitive differentiator; instead, the company disclosed the MMX instruction set in full detail nine months before it shipped the first processor to implement it, enabling AMD and Cyrix to ship compatible processors just a few months after Intel.

The first Pentium/MMX chips had barely been shipped, however, when rumors began circulating about the next round of extensions, called MMX2, which Intel still has not acknowledged. The prevalent rumor, however, is that MMX2 will do for single-precision floating-point what MMX does for integer calculations: enable multiple operations to proceed in parallel, accelerating 3D geometry calculations. We expect MMX2 to first appear in Katmai, an enhanced Pentium II family member due in early 1999.

The opportunity to create an FP equivalent of MMX, as well as to provide other instructions to speed 3D geometry, has not gone unnoticed by Intel's competitors. The x86's FP architecture is so weak that it isn't hard to implement something far better. Unfortunately, AMD, Cyrix, and IDT are following mutually incompatible paths.

Much of the justification for the proprietary extensions comes from a belief that they can be made useful simply by optimizing Microsoft's Direct3D code, using the extensions for geometry and lighting calculations. Applications (primarily games) that use Direct3D for lighting and geometry would then be sped up without changes to the applications.

Microsoft appears to be open to working with the chip vendors, but they will have to get working systems into Microsoft's hands well before the planned 2Q98 ship date for DirectX 6 if support for their chips is to be in this release—and there are no additional releases planned until 1999. Cyrix plans to support its extensions through its own 3D libraries, possibly focusing on OpenGL instead of DirectX.

Unfortunately, virtually no games today use Direct 3D for lighting or geometry, although many use it for rendering. DirectX 6 is supposed to be much faster at these functions, making it more attractive for game developers to use them.

Recognizing that the shift to Direct3D geometry may be slow, AMD and Cyrix also hope to convince game developers to rewrite their proprietary geometry and lighting code to use their extensions. (Centaur does not believe developers should use the new instructions directly and is depending entirely on DirectX support.) AMD and Cyrix may get some support for this—especially if they provide technical help and financial incentives—but many game developers will simply wait for MMX2, or avoid the entire issue by moving to Direct3D geometry. How much work can they justify doing for a short-term performance boost for a small fraction of the user base? How many variations of their code will game developers be willing to develop and support? They might support one of the new extensions to gain a "hot box" demo platform, but not all three.

The new extensions are bad news for 3D chip vendors planning to offer Direct3D geometry acceleration in hardware. The geometry speedups provided by the software extensions will make host-based geometry harder to beat—and in any case, Microsoft's current plan is not to support hardware geometry acceleration in DirectX 6.

Once processors with MMX2 ship, Intel's competitors are likely to revise their processors to be compatible with Intel's, making their proprietary extensions largely irrelevant. If their proprietary extensions provide superior performance to MMX2, if they can get Microsoft to include code that uses their extensions in future versions of DirectX, and if game developers use Direct3D geometry and lighting, the extensions could continue to deliver benefit even if developers don't use them directly.

The competitive processor marketplace encourages innovation, as each vendor seeks a competitive advantage—but much of this innovation may be for naught. Intel did the industry a service by allowing MMX to be standardized, but now Intel's competitors want to move more quickly than Intel. Had Intel laid out a roadmap to MMX2, or moved more quickly to implement it, the fragmentation of the architecture could have been avoided. Intel seems willing, at this point, to let the architecture fragment to weaken its competitors. For Intel's competitors each to go its own way, however, is a serious mistake: it is time for them to work together and support a common standard if they are to succeed in establishing non-Intel extensions. M

See www.MDRonline.com/slater/x86frag *for more on this subject. I welcome your feedback at* mslater@mdr.zd.com.