Microprocessors: The Next Five Years

By Michael Slater, for EE Times, October 1996

Now 25 years old, the microprocessor is still in an early stage of its evolution. During the next 25 years, huge changes are likely. But if we narrow our focus to the next five years, the entrenched positions of today's microprocessor makers and the applications already identified become the dominant forces, and we can predict the future with somewhat more accuracy.

Digital consumer electronics is one of the most exciting new areas for embedded microprocessors. These applications are taking the volume lead from traditional 32-bit embedded applications, such as laser printers and network routers, and causing shifts in market share among architectures. Hitachi's SH family has shot to first place among 32-bit embedded RISCs, driven largely by the success of Sega's video games. The MIPS architecture remains a strong number-two, with video games from Sony and Nintendo expanding its market beyond printers and other control applications. Handheld organizers, formerly the province of limited processors and proprietary software, will get a big push from Microsoft's Windows CE, which will run on Hitachi's SH, MIPS, ARM, and others. DVD players will replace VCRs will special-purpose video computers, offering new opportunities for fast but anonymous processors.

As available transistor counts skyrocket and cost pressures increase, integration of system functions along with processor cores will become routine for high-volume embedded applications. For this segment of the market, design tools and cell libraries could be more important than instruction sets.

In the desktop computer market, the overwhelming dominance of Windows has led to the equally overwhelming dominance of the x86 architecture. Of the five leading RISCs, Hewlett-Packard has already thrown in the towel and decided to join forces with Intel on the future IA-64 processors. SPARC and MIPS will continue to be used in systems from Sun and Silicon Graphics, respectively, as well as by some of their close partners, but their opportunity to become more widespread in the mainstream desktop market has passed. IBM and Motorola have the highest-volume desktop RISC processor with PowerPC, but its only volume market today is Macintosh—and this isn't likely to change for at least two more years. Digital continues to push Alpha toward the mainstream market, but this will be a long, slow process.

For the next three years, then, there is little question that the x86 architecture will remain dominant. Intel's x86 competitors are likely to capture a minority share of the x86 market, focused toward the lower-priced products, while Intel holds the majority of the units and the great majority of the revenue.

Within five years, however, things could begin to shift—but Intel's position still looks strong. A descendent of Windows NT will replace Windows 95, eliminating the x86-only operating system from Microsoft’s lineup. Intel will have its Merced processor, the first implementation of the IA-64 architecture, so even Intel will be encouraging the software community to focus on portability. Until Merced ships, Intel can probably keep the vast majority of the PC market happy with enhanced P6-family processors—and once Merced does ship, it may be difficult for the RISCs to maintain a significant performance lead. It is far too early to call this battle; Intel will have the advantage of a massive manufacturing capability and, presumably, the best x86 emulation, while Alpha, PowerPC, and possibly others will seek to offer better performance and a way out of an Intel-dominated world.

One wildcard is the emergence of Java as a processor-independent form for distributing software. Should a sizable base of Java applications appear, this could reduce the barriers between architectures.

To continue doubling performance every 18 months, processor designers will have to resort to ever more exotic techniques. Memory bandwidth is a key limitation, and many processors will stay on the performance curve for the next year or more by adding larger on-chip caches, fast level-two cache buses, and high-performance memory interfaces. More execution units will also be added to the processor cores, but with today’s software base, the point of diminishing returns has already been reached: additional execution units tend to sit idle for lack of adequate parallelism in the code. In the long run, new compiler technology—and possibly new instruction sets—will be needed to take advantage of the parallelism that 100-million-transistor chips will be able to implement. Another course, sure to be taken by some vendors, it to implement a coarser level of parallelism by integrating multiple processors on a single chip.