THE EDITOR'S VIEW

Limits to Integration

By Michael Slater

The phenomenal increase in transistor density made possible by ever-improving semiconductor technology has made it possible to integrate complete systems on a single chip. With million-transistor chips having become commonplace and three-million transistor chips now going into production, it is possible to include not only a powerful CPU core but also an assortment of peripheral circuits, interface logic, memory, and other supporting functions on the same chip.

Despite the capability to fabricate chips that include a variety of functions in addition to the processor, the pace at which this is being done is slow. Among the factors fighting the integration trend are the cost of high-pin-count packages and the decrease in yield as chips get larger. A key problem, however, is the difficulty of deciding what to put on the chip and then designing all the required functional blocks. Exacerbating this problem is the fact that chip designers are now forced to make system-level design decisions, and they are often ill-equipped to do so.

In the embedded processor market, the desire to keep the range of applications as broad as possible limits the amount of functional integration that manufacturers are willing to do. Embedded processors go into a wide range of applications, each with somewhat different needs, and the market (at least for high-performance processors) is small enough that vendors don't want to fragment it further than necessary.

There have been some exceptions for the highest-volume application areas. For example, LSI Logic's LR33020 is designed for X terminals, and AMD's 29200 is designed for laser printers.

In many cases, it seems that useful functions are often left off simply because the semiconductor maker can't decide what to include. Relatively few high-end embedded processors include programmable chip select outputs, for example, yet virtually all applications would benefit from them and the amount of silicon area required is small. Pin count is one limiting factor, but it is not the only one.

When we asked one maker of a high-end embedded processor why they didn't include programmable chip selects, they replied that when they asked prospective customers what features were important to them, programmable chip selects weren't requested. This may be an example of listening to customers too much—it is the chip maker's job to innovate. Customers often don't know what is possible, and if vendors just build what customers ask for, progress will be slow.

The issues in the general-purpose microprocessor market are slightly different. Here we have one application with enormous volume, but there are still many variations—different display standards, memory configurations, I/O buses, and so forth, for different price points. So far, the only fully integrated PC chips are C&T's PC/Chip and Vadem's new VG-230 (see p. 4), and both are based on 8086-type cores.

Considering only transistor count, there is no reason why a 386-core single-chip PC (less memory) cannot be built today; even the 1.2-million transistor budget represented by the 486 would be adequate, and a state-of-the-art 3-million-transistor chip would provide room for a sophisticated set of peripherals.

Eventually, highly integrated processor/system logic chips are likely to dominate the mainstream portable and desktop computer markets, but this transition will occur slowly. Pin count is a key issue here: a fully integrated, high-end PC chip would require about 500 pins for a full set of traditional interfaces. Another key problem is assembling the design expertise for all the functions required—display controller, disk controller, and so forth. The number of different functions also increases the chance that something will go wrong, making debugging a major concern.

In the embedded arena, ASICs combining a standard CPU core with a mix of standard and customer-specific peripherals could become popular. Until now, design tools and testing issues have slowed progress in this area, but these problems could be overcome. LSI Logic is making a major push in this direction with its CoreWare program, which offers both MIPS and SPARC CPU cores. A similar approach will eventually appear in the PC arena. Already, NEC offers this capability with its V20 and V30 cores; this is how Vadem was able to craft its VG-230. Eventually, TI is likely to offer the Cyrix processor designs as megacells within its ASIC library, and companies such as LSI Logic must surely be on the search for 386-compatible cores.

In the final analysis, however, it may turn out that there is often little advantage to putting everything on one chip. For many functions, there is no performance gain to putting them on the same chip as the processor. A system partitioning with a standard microprocessor chip and an application-specific "everything else" chip may often be the most effective solution.

Whatever approach turns out to be the most successful, one thing is clear: the enormous increase in transistor budgets over the coming decade will continue to provide designers with more transistors than they can figure out how to use—a glorious problem to have. ♦