Instruction Sets Not Key to Performance

Other Performance Factors Outweigh Need for New Architectures

New instruction sets are old hat. A couple of years ago, a few processor vendors were looking at new instruction sets, mainly in response to Intel's IA-64 initiative. Building an instruction set from scratch is fun—and always a good PhD dissertation—but in the end probably isn't the best way to gain performance. Realizing this, these same CPU vendors are now busy designing new memory hierarchies instead.

These vendors have concluded that a new instruction set architecture (ISA) just isn't worth the trouble. There's nothing stopping IBM or Sun or Compaq from designing a new ISA. Compatibility with the installed base is always a concern, but the same hardware and software translation strategies that will be used for Intel's IA-64 could be used for other architectures as well. But putting together a new ISA, test suites, development tools, and compatibility tools is an enormous task for which the payback appears small.

True, a new ISA can take advantage of new research in computer architecture. The past decade has seen important advances in branch prediction, predication, speculation, SIMD execution, and cache design, among other areas. A new ISA can incorporate any new feature its creators desire. Existing instruction sets, however, can be (and have been) enhanced as well. In fact, vendors have added all of the above features, in some form, to existing instruction sets.

For example, Intel touts predication and speculation as key advantages of IA-64. But several RISC vendors have speculative loads or prefetch instructions that offer many of the same benefits. Similarly, these RISC architectures have conditional-move instructions that are much simpler than full predication yet eliminate nearly the same number of branches, although not quite as efficiently.

Other features in IA-64, such as the larger register file, are difficult or impossible to retrofit onto existing architectures. Some applications will undoubtedly benefit from these new features. The ultimate question is how big this benefit will be. I doubt that it will be more than 20–30%.

Let's look at history. The big ISA battle has been RISC versus CISC. This debate has died out in recent years as Intel demonstrated that none of the problems of CISC were insurmountable. On SPECint95, Intel's CISC processors currently outscore all RISC processors except for the Alpha chips. If instruction sets made much of a difference, the other RISCs would all be faster than Intel's P6.

On SPECfp95, the P6 does in fact trail all the high-end RISCs. This performance shortfall is due more to the P6's weak FPU than to its convoluted floating-point instruction set. The P6 takes two cycles to start one FP multiply and one FP add; the best RISC processors can launch four times as many FP operations. If Intel ever included a competitive FPU in an x86 processor, that chip would probably approach the performance of contemporary RISC products.

Let's look at the future. New ISAs are being designed to deliver more instruction-level parallelism (ILP) than existing architectures, which are not well designed for parallel execution. But as my colleague Keith Diefendorff points out (see MPR 12/28/98, p. 3), the amount of ILP in existing applications is limited; today's four-issue processors wring nearly all of the ILP from the majority of applications.

Even if future compilers can extract more ILP from applications, the bigger problem facing processor designers is keeping ever more powerful CPUs from data starvation. On most applications today, processors spend more cycles waiting on memory than they do processing instructions. New innovations that improve CPU efficiency will merely increase the number of memory stall cycles.

So why are Intel and HP developing a new instruction set? The reasons have little to do with technology. Intel needs a tool to attract support in the high-end workstation and server markets, and ultimately the new ISA will help Intel distinguish itself from AMD et al in the PC market. Both HP and Silicon Graphics grew tired of the financial burden of supporting their own architectures and decided to let Intel spend some of its billions instead.

The biggest factor affecting performance in the future, as in the past, is the skill of the design teams in choosing the best implementations. Using nearly identical ISAs and similar process technology, RISC processors today span a 2:1 range in performance, far greater than any benefit likely to come from a new instruction set.

The most important questions concern implementations. Can Intel deliver an IA-64 processor that meets its schedule and performance goals? Can Sun overcome its history of laggard uniprocessor performance? Can IBM deliver processors with competitive clock speeds? Will Compaq continue to invest in an Alpha architecture with a small installed base? Will Intel be able to match the exotic memory architectures of its competitors? These and similar issues, not brand-new instruction sets, will determine the performance leaders over the next five years. M