Economist.com SCIENCE & TECHNOLOGY

Chip design

Supercalifragilisticexpialidocious

AS NEW materials come to hand, architects find new uses for them. With cheap steel beams and glass came the modern skyscraper, complete with the new technologies of high-speed lifts and service cores. The same principle holds even if the scale of the building is smaller -100m times smaller, say. Computer architects make the most of the building components available, and the types of chips that result all need different technologies to work.

The changes in the materials available to chip makers are changes in both abundance and cost. The number of components that can be put on to a chip has grown by 40% a year. The cleverness of designers, who have used the capacity that comes with more power ever more effectively, has parlayed this into a 50% annual increase in performance. The most striking change of all has been from old-fashioned CISC architecture to new-fangled RISC chips. Now researchers are looking at another big change that could help them to get the most out of the next generation of even more tightly-packed chips: very long instruction word (VLIW) sets.

The earliest microprocessors used complex instruction set computing (CISC). The main constraint was that the memory in which programs are stored, random-access memory RAM, was expensive and slow. An alternative, read-only memory (R0M), was both cheaper and faster; but it could not be changed. In order to minimise the tedious business of getting into a RAM, the instructions within it were kept quite complex; they were like German compound nouns that have lots of subdivisions. They were then broken down into simpler commands stored in a ROM.

Around the beginning of the 1980s the price of RAM5 crashed and their speed increased; visiting a RAM was no longer something to be minimised. So academics at Stanford and Berkeley Universities suggested simplifying the instructions and saving time on the business of breaking them down: this was reduced instruction set computing (RISC). Programs would visit a RAM much more often and spend less time turning what they found there into simpler instructions. The computer could get its commands in a steady smooth stream, rather than in compound job-lots.

It took a lot of work to turn this idea into new chip architectures, but that has now happened. For high-performance tasks, RISC has shown marked advantages over CISC. The older technology is now hanging on only because of its vast installed base. RISC chips and CISC chips are not interchangeable; to get the best out of either takes a computer designed only with that chip in mind. Most families of computers today were designed for CISC. But RI5C is making inroads in markets where performance is highly valued. It is even coming into desktop PCs, thanks to the PowerPC chip made by IBM and Motorola.

The building materials, however, continue to get cheaper. You can now have several million transistors on a chip—more than a single processor needs. One use for this extra capacity is to put some of the RAM that used to be kept on a separate chip into the microprocessor itself, making it even more accessible. Another is to have different parts of the chip operating independently, in parallel. Chip designers have begun to duplicate some of a chip's abilities, so that it can do things twice as fast, sending one bit of the program to one set of circuits and another bit elsewhere.

This “instruction level parallelism” lets modern RISC chips execute two to four instructions at a time. The cost, though, is that the incoming instructions have to be sorted very carefully to make sure that calculations that depend on each other are not farmed out to different circuits that do not interconnect. The logic needed to do this vastly complicates the design effort and runs against the RISC designer's first commandment:
“Keep it simple, stupid”.

The VLIW approach is meant to foster further parallelism, allowing the 10m-20m components that will be available by the end of the decade to do useful work in perhaps 30 separate processing centres spread across the chip. The secret lies in removing all the fuss of splitting the jobs up from the chip's ambit. Instead, that job is done with software: specifically, with some software called a compiler. This turns the sort of language that computer programmers use— such as C or Fortran—into the simple ons and offs of a computer's own instructions. In VLIW computing the compiler sorts the program into chunks that can be executed in parallel. These chunks (the very long words) are then easily dealt with by individual bits of the chip.

When VLIW was first developed it had the disadvantage that the programs needed to be put together with the specific machine that would run them in mind. That was a drawback which could have restricted the technology's potential market. Hewlett-Packard now claims it has solved this problem. Bob Rau at Hewlett-Packard, who was one of the initial developers of VLIW, thinks that the technology's chances are further improved by the growing importance of multimedia computing, which requires computers to surge through reams of video and audio data. Once it seemed that only scientists and engineers would need VLIW's power; now all sorts of people will.

Hewlett-Packard is now working with Intel on the technology, in a collaboration to which both parties expect to devote $1 billion. Even at this price, they hope to get a bargain. Intel, which makes the CISC chips that dominate the PC market, has been good at low cost manufacturing and aggressive marketing. Now it could find the best way of making the next generation of chips, too.

THE ECONOMIST FEBRUARY 4TH 1995, p. 75