

Transputers

Inmos puts its head in the lion's mouth

Tokyo

SELLING microchips to Japan is not a task for the faint of heart. But Inmos, the United Kingdom's only microchip company and a pygmy compared with the eight giant Japanese electronic corporations, is trying to do just that. Last week the revolutionary product on which they pin their hopes — the transputer — was launched in Tokyo, as well as in the United States and the United Kingdom.

To avoid being crushed underfoot, Inmos has aimed at a product that does not compete directly with anyone else's. The transputer is claimed to be the first "computer on a chip"; on less than a centimetre square of silicon are a programmable processor, a 2K random access memory and a special set of four two-way communication channels. These features make the transputer not only a powerful microprocessor in its own right but also give it the capability to be linked up to many

other transputers — tens, hundreds or perhaps even thousands. Thus the transputer may provide a building block for computers of the future that work in parallel rather than sequentially.

There are good reasons for building parallel computers, although enormous conceptual difficulties remain in building a general-purpose parallel machine (see below). In principle, each time the number of processors is doubled, the processing time is halved — and the sky is the limit. Of course, it is not so easy to divide up problems among processors to get linear increases in speed. But Inmos does have a few convincing demonstrations of what can be done, given the right problem. One example is that of generating a two-dimensional picture, complete with shadows and highlights, of several three-dimensional coloured objects illuminated by lights. The problem is not easy because each object will contain reflections of all

the other objects as well as the lights. It can be solved by tracing the path of each ray of light; and as the rays are unaffected by one another it is possible to divide the work among processors — the more transputers Inmos puts onto the problem the faster it is solved. Similar logic has been applied in building a machine to encode and match fingerprints for Scotland Yard.

The possibilities for inventive designs incorporating transputers are endless. And it is those who are looking for new parallel solutions to tricky problems who are likely to be among the first customers for the transputer. Indeed, such people are already banging on Inmos's door. But for financial success Inmos has to look beyond these customers to the conventional microprocessor market and that will be a harder nut to crack. Used as a single 32-bit microprocessor (the kind of chip that forms the heart of personal computers, for example), Inmos claims that the transputer is still very considerably faster and more efficient than microprocessors produced by Intel or Motorola, the two US companies whose designs dominate world markets. And it is very easy to add more transputers when more facilities are required. But Intel and Motorola are already well established.

Beyond the conventional microprocessor market may come a third wave of buyers as the possibilities offered by parallel processing become more widely appreciated by conventional programmers. Indeed, it is even possible that OCCAM, the transputer's special programming language, will become a world standard for parallel applications.

But Inmos has first to establish itself in the microprocessor market. The next two years should show whether the transputer will be a world-beater or whether, like the rear-engined jet and the hovercraft, it will turn out to be one of those British ideas that are too far ahead of their time.

Alun Anderson

Why should computers go parallel?

Tokyo

A CONVENTIONAL computer operates step-by-step by fetching data from a central memory store, operating in it and returning the result to the store. To crack huge problems, like those of modelling the behaviour of the atmosphere or nuclear reactions, or to produce the computational power needed for "intelligent" applications, computers are being built with central processors that can perform one operation after another at ever higher speeds. But the limits to this approach are already on the horizon.

The day is coming when computers will perform operations so quickly that even at the speed of light there will not be time to fetch data from memory stored more than a few tens of centimetres away. So computers will have to become more and more compact, but as they do it becomes impossible to disperse the heat they produce. That is one reason why researchers are trying to find ways to use many processors working in parallel; another is that the staggering engineering problems that have to be overcome to build the next generation of supercomputers will make them prohibitively expensive.

Solving a problem in parallel is much like tackling a real-life organizational problem, say that of running a journal. Different parts of the total task can be assigned to different editors, all working in parallel. But their efforts have to be synchronized and information passed back and forth between them (and outside con-

tributors) in order that the final goal, production of a complete journal, can be achieved. The difficulties of dividing up a problem, ensuring that different parts are done at the right time, passing the right information back and forth between processors, and making sure no processors are sitting idle, have to be solved to build an efficient parallel machine. Some types of problem will require individual processors to be powerful and fairly independent, as in image processing where the same operations can be performed simultaneously on every part of the image. Others will require enormous information exchange between different processors and thus enormous interconnectivity.

The transputer can provide the basic building block for the construction of such parallel machines, as its on-chip memory and special communication channels enable it to talk to other transputers while performing operations. Even more important is that the transputer has been built around a special language, OCCAM, developed at the University of Oxford for parallel applications. OCCAM commands make it easy to specify what set of processes are to be performed in parallel and to define communication channels which come into use when receiver and sender are ready to pass information. Indeed, OCCAM is the first practical language to provide a simple means of controlling concurrent operations; it already has 5,000 users worldwide including an enthusiastic band in Japan.

Alun Anderson

Soviet hostages dead

Two Soviet geologists who were kidnapped by rebel forces in Mozambique in August 1983 are dead, the Mozambican foreign ministry has announced. The geologists, Yurii Gavrilov and Viktor Estamin, were in Mozambique under a scientific and economic cooperation agreement, as advisers to the Mozambican mining industry in Morrua, Zambezi province. During a rebel raid on 21 August 1983, 24 Soviet geologists were captured; most, however, were subsequently released as a result of military action by government forces and diplomatic negotiations. Gavrilov and Estamin remained in rebel hands.

This brings the total of Soviet victims of the Morrua raid to six; two were killed in the raid itself, while two others were known to have died of hardship during their captivity.

Vera Rich