Japanese computers

Fifth-generation machines here

Tokyo

JAPAN'S fifth-generation computer project has taken a big step towards its preliminary goals with the successful construction of a prototype relational database machine — the first of its kind.

The fifth-generation computer project aims at the development, by the 1990s, of a completely new form of computer that can communicate in everyday language, "reason intelligently" and bring enormous amounts of stored knowledge to bear on any problem that the user might set it. The machines are intended to be cheap and reliable enough to be used in homes and offices. Some talk as if they will revolutionize society by making freely available a huge fund of expertise that will vastly "amplify human intelligence".

When the project was announced by the Ministry of International Trade and Industry (MITI) in 1981, its sheer boldness and originality astonished people in the West. Japan, considered weak in fundamentally new ideas, seemed overnight to have produced a plan that would give it an unassailable lead in advanced computer technology. One expert even uncharitably described the plan as a "technological *Mein Kampf*".

Since then, Western governments have responded more rationally with new projects of their own. There are, for example, advanced information technology projects in the United Kingdom, France and West Germany, together with the European Communities' Esprit programme, while the schemes of the United States' Defense Advanced Research Projects Agency were directly inspired by Japan.

The fifth-generation computer will be composed of three main functions: an intelligent interface allowing it to communicate flexibly with humans by speaking, writing or drawing; problem-solving and inference functions allowing problems to be solved by deductive and inductive inference as well as by making sensible guesses where knowledge is incomplete; and a knowledge base function able systematically to store and retrieve not only vast quantities of data but also judgements and test results.

It is within this third function that 45 researchers at the Institute for New Generation Computer Technology (ICOT), set up by MITI to run the project in cooperation with Japan's big eight electronic manufacturers, have now announced a major step forward. Success with the construction of the relational database machine follows by five months the completion of the hardware of one of the other two essential parts of a prototype fifth-generation machine — a sequential inference machine.

The novelty in ICOT's machine is that relational software has been replaced by specially developed hardware. Not all the control software is complete so the machine is not actually doing anything yet. But when in operation, it will be the first true relational database "machine", working two orders of magnitude faster than anything else available. Named Delta, the device consists of two parts: a relational processor complex containing pipelined relational algebra engines (built by Toshiba) and a hierarchical memory subsystem (built by Hitachi), plus an interface (built by Oki). The aim is to be able to connect it with the prototype sequential inference machine by next spring. ICOT hopes then to have two-thirds of a prototype fifth-generation machine, providing the programming environment in which to tackle the really big problems ahead.

Connecting the two together requires that both operate on the chosen "kernel" language of extended PROLOG, the advanced logic programming language invented in France and developed in the United Kingdom. There are several tricky software problems to be solved.

That both ICOT's sequential inference machine, which broke new ground as hardware by allowing extended PROLOG to be used as its machine language (see below), and the relational database centre on relations between things is no accident. ICOT's visionary director, Kazuhiro Fuchi, sees the predicate logic they use, established by Godel in the 1930s, as the key to building knowledge information processing systems. Indeed, he has said that the present generation of computers, based on the Turing theory published in 1936, may one day be seen as a "gigantic historical detour".

The fifth-generation computer project is scheduled to last ten years. The first three years, up to next spring, seem likely to achieve the goal of a prototype machine to be used as a base for further research. Eventually, however, machines will have to work at incomparably greater speeds than at present. The goal is 100 to 1,000 million logical inferences per second (Lips), compared with just 30,000 Lips for the present machine.

This speed can be achieved only by switching to some sort of parallel architecture for the sequential inference machine, probably based on a data flow model (in which the University of Manchester is an acknowledged leader). The project also has to solve the stupendous problems of developing syntactical analysis systems allowing communication with humans in natural language, making a complete revision of PROLOG for parallel use, developing knowledge-base machines utilizing parallel operations and pattern information processing systems.

What has been achieved so far consists of the rapid and determined development of ideas well researched elsewhere. The four-year intermediate stage of the project beginning in 1985, however, requires radical innovations — on time. Those years will decide whether a Japanese fifth-generation computer will emerge in the 1990s or whether it will have to be postponed until the world of information technology is much further advanced.

This year's budget for the project is Y5,100 million (£16 million). Later budgets, it is said, will depend upon "progress in research". **Alun Anderson**

Logic and language for the new generation

RELATIONAL databases are not in themselves new — their mathematical basis was laid out in 1972. What makes them distinct from more conventional data bases, and particularly suits them to "intelligent" applications, is their simplicity of organization. Entries consist purely of relations (essentially two-dimensional tables).

This provides access to any datum in the base either directly or through association and allows data sets to be rearranged in new combinations. If, for example, a set of statistics about university students and their examination results is stored in a relational data base, it is easy to ask the base to provide, say, the names of all students who scored exactly 85 per cent in their examinations. Traditionally, this would be accomplished only by searching through the whole data base. The facility to retrieve data flexibly (by relational algebra) has obvious advantages for intelligent operations, for it will never be clear in advance just how knowledge will be used.

PROLOG, the chosen kernel language, was originally devised for solving problems that involve symbolic representations of objects and relationships between them. To the average programmer, accustomed to writing sequential instructions, PROLOG is something of a surprise.

Thus programmes are written in terms of facts, for example:

```
Tim likes doughnuts
```
written as
```
likes (Tim, doughnuts)
```
rules, for example:
```
Tim likes any object
provided it is sweet
```
and questions, for example:
```
Does Tim like doughnuts?
```
written as:
```
? likes (Tim, doughnuts)
```
which will cause a search to be made to see if a precisely matching declaration can be found in the data. Building from these very simple roots, extremely complex logical arguments can be advanced and complicated inferences be made from one statement to another. Alun Anderson