Fifth generation computing 1984 : a report to the Ministry of State for Science and Technology, Government of Canada

Fifth Generation Computing

1984:

A Report to the Ministry of State for Science & Technology Government of Canada

by

Romas Aleliunas
Forbes J. Burkowski
Randy Goebel
Eric G. Manning

Department of Computer Science
and
Institute for Computer Research
University of Waterloo

7 December 1984

This document reports on a two-week visit to Japan during October and November, 1984. The authors wish to acknowledge the contribution to this report made by Dr Douglas Dyment of Sagan Technologies, Vancouver, and by Dr W. Morven Gentleman of the National Research Council of Canada, both former professors of the University of Waterloo who accompanied us to Tokyo. Dr Dyment and Dr Gentleman made numerous valuable observations and contributed to the content of this report in several ways. We also wish to thank the Department of External Affairs and the Ministry of State for Science & Technology, for financial and moral support. We are especially indebted to Messrs. J. MacDowall and D. Silverman of MOSST, and to Mr. Peter Eggleton of the Embassy in Tokyo. Finally, we are indebted to Dr. Yutaka Matsushita of OKI Electric Industry Company Ltd. (OKI-Denki), who arranged several of our visits and provided invaluable introductions. Nonetheless, responsibility for errors or omissions rests with the authors.

RA
FJB
RG
EGM

Waterloo, Ontario
7 December 1984

Contents

Executive Summary

• We believe that the industrialized economies will increasingly depend on **information technology (IT)** for their wellbeing. All will depend increasingly on IT to improve the competitiveness of their traditional economic mainstays. In addition, some (including both the US and Japan) will earn more and more of their national incomes from the sale of information-based products and services. In either case, the effective use of IT will be critical to future wellbeing.

• Moreover, recent trends in IT are leading us to systems which are "intelligent" in their use of data (for example, able to make logical inferences from given "facts"), and which are much easier for non-experts to use than most current products. Virtually every field of human social and economic activity will be altered by such systems; the first major program aimed at developing them was the Japanese Fifth Generation Program, announced by MITI in 1982.

•These trends will require us to develop information processing systems along the general lines pursued by Japan's Fifth Generation program: machines capable of making **logical inferences** quickly and inexpensively, databases capable of storing both information and rules governing the use of that information (so-called **knowledge bases**), and user interfaces based on **speech and images** as well as text. In short, computers which can be treated more as **human assistants** than as moronic machines.

• The building of such machines is mostly a matter of **systems integration** - the pulling together of achievements from several areas of technology. The key areas of technology which underpin the Fifth Generation are **microelectronics** – especially Computer-Aided

Design of ULSI chips, **computer architecture** - especially distributed & parallel architectures, **software engineering**, and **artificial intelligence**. Any worthwhile R&D program must emphasize these areas; the US **Strategic Computing Program**, the UK **Alvey Programme**, the EEC **Esprit Programme** and of course the Japanese (Institute for New Generation Computer Technology or **ICOT**) program, all emphasize them heavily. Heavy emphasis on **industry-university** collaboration is the other essential feature of a successful program.

• While there is plenty of room for argument about the details of the Japanese approach (e.g. the use of Prolog rather than some other language), it is clear that they are pursuing a focussed, cohesive, well-financed R&D program aimed at the right strategic objectives. No single result from Japan to date is technologically breathtaking; much of their work is catch-up and some is misguided or naive. **However**, the objectives of the program's first Phase - prototype and tool development - have been achieved, and achieved on schedule. (We remark that the Japanese made few single steps that were technologically breathtaking, in the consumer electronics, machine tool, shipbuilding, optical or automobile industries, either.)

• No-one knows what the market for knowledge machines might be in five years' time. It could be substantial, perhaps comparable in size to the market for "super mini" computers and thus able to sustain a few good-sized companies. If the statements made above turn out to be accurate, the world market for **knowledge processing** (hardware, software and data) could be enormous in ten years' time, comparable in size to today's total market for Information Technology. In this event the Japanese will be in the driver's seat, as they are training substantial numbers of engineers in the new technologies, and are gaining invaluable experience through the construction and tryout of prototypes and the conduct of basic research. They have also succeeded in setting an **agenda** to which the rest of the world feels compelled to respond - the ICOT Conference is the premier forum for the reporting of progress, and advocates in other nations defend and justify their research proposals by comparison to the Japanese program.

• Most emphatically, the Japanese do **not** yet have such a commanding lead, that it is pointless for other nations to enter the race. Indeed, large portions of the potential market are not being addressed by the Japanese. Other parts are addressed badly; still others are of such importance to Canada that we should address them in any event. Here are some promising areas which Canadians could pursue:

- sets of **tools for building software** (software development environments) using knowledge-based technology. If the Japanese reach their stated goal of building such environments and if no comparable effort occurs in Canada, our current lead in software development will be wiped out.

- Japanese activity in **natural language translation** has increased markedly in the past six months; about six new products have been introduced to the Japanese market. The work is based on the research of the ICOT program, and its importance to a bilingual nation is clear. The extreme difficulties of NL translation should however be kept in mind, and modest but achievable goals should be set.

- The construction of an **expert system** for a particular problem is far from a matter of routine engineering practice. However, more than enough successes have been chalked up, and the potential payoffs are impressive enough to justify a substantial Canadian R&D effort. Research on better tools for expert system construction (logic programming, knowledge representation, software engineering, hardware architecture,...) is required. Also, specific systems to improve efficiency in traditional Canadian industries - insurance companies & banks (product selection), natural resources (interpretation of oilwell data, selection of forest stands), transportation (automated travel agents), medical care, and education (Computer-Aided Learning based on natural language front ends and expert system techniques) - are attractive. Discussions should be held with the industries involved to validate or modify this sample listing.

- Canada has outstanding researchers working on **computer vision**; applications to robotics and the processing of images from satellites are just two of many exciting possibilities here.

- Although the Japanese have developed prototypes of small, eventually desk-top **computer workstations** to support knowledge processing, we feel that Canadian efforts in this area could be rewarding. Such machines may have a very large market in five to ten years, as they would be attractive for use with expert systems of all kinds, natural language translators, natural language front ends (to interface humans with traditional computer systems via natural language discourse), and knowledge database systems. Moreover, the current Japanese prototypes are of uninspired design and contain at least one serious design misjudgment.

We recommend the immediate creation of a joint industry-university research program, funded by government and industry and administered by the Natural Sciences and Engineering Research Council, to strengthen Canadian research in the disciplines underlying Fifth Generation computing and to support the pursuit of these areas by Canadian industry.

• In a previous trip report (Manning & Wright, 1983) the general level of Japanese computer expertise was assessed. The main finding was that Japanese expertise tends to lessen as the topic becomes more abstract. Thus Japanese packaging, mechanical design and production engineering of computer hardware are superb. However, their computer **architecture** (high-level hardware design) and **software** design are at best imitative of the best in the West, and at worst are just plain terrible. In fact, Japanese software designers appeared to be not fully aware of Western advances in these fields.

Our findings from this trip are similar, with one notable exception. A few of the software designers at Mitsubishi's Kamakura lab were fully up to speed on Western technology. In addition, their healthy criticism of previous work and the obvious pleasure which they took in creation suggest that competitive software may soon emerge from Japan. Interestingly, they were working on the basic control software for the Japanese prototype workstation for knowledge processing - the Sequential Inference Machine- Personal or SIM-P.

Visits: the Mitsubishi Lab at Kamakura

We visited this lab on 31 October. Mitsubishi is not yet well-known in the West as a computer manufacturer; nonetheless they have a Computer Division with three Works and a Laboratory.

The Lab is new; it was started in 1981. It has about 300 professionals in five departments, namely:

- computer & software development
- system development
- information transmission
- optics & microwaves
- optical information systems.

The General Manager is Dr E YAMAZAKI.

It is remarkable that such a large Laboratory could be set up and brought to a state of high effectiveness in such a short period.

The Computer & Software Department of the Laboratory has about 80 professionals working on several topics. The hardware group developed the prototype of the SIM-P knowledge workstation for ICOT (working from ICOT's high-level design), plus a highly parallel superspeed computer intended for scientific & engineering work, plus a machine with both the IBM 370 instruction set and the instruction set of the former SDS Sigma series. This last machine is alleged to be Mitsubushi's next major product offering.

There were also groups in the Computer & Software Department working on:
Computer Aided Design for VLSI (software tools for microchip design, test generation and simulation), **software technology** with a heavy emphasis on Fifth Generation topics such as databases and natural language processing for Office Automation, **software engineering** (they sell a CAD system to support software engineering), **signal processing** and **process control.**

The System Development Department has yet another workstation or large personal computer based on the Motorola 68000 microprocessor and the UNIX operating system. They also have several interesting products for local area networking, including a frame-formatted network or LAN running on optical fibre at 32 Million bits/sec (pretty fast), a CSMA-CD LAN running on fibre at 10 Mbit/sec, a star-topology PBX which carries both voice and data, and an optical-fibre crossbar switch for 32 ports. This set of products will permit Mitsubishi to offer most of the interesting combinations of LAN technology for the Office Automation marketplace.

The Information Transmission Department is working on pattern recognition for Kanji (Chinese) characters, images and speech; imaging techniques for facsimile transmission of documents, coding, and vertically-encoded disc drives. (The latter will increase the capacity of magnetic discs by a factor of 10 if it can be proven in.) We heard very little about the Optics & Microwaves and Optical Systems departments.

Members of our team made the following comments about the Mitsubishi Lab:

- highly apppropriate selection of projects

- clear understanding of the competition

- impressive level and range of activity for the age and size of the Computer & Software Department

- world standard in terms of expertise, but nothing earthshakingly original

- well-equipped to take ideas from elsewhere and turn them rapidly into products

- to the extent that the market for knowledge machines develops, they'll have a big lead on the competition. (They will have delivered 26 SIM-Ps to ICOT by January 1985.)

- intelligent and knowledgeable people exuding an air of excitement and pleasure in creative work.

- knowledge-based computing sure to be extremely important in the future. Mitsubishi is sure to have big edge, even if Prolog doesn't become the favoured language - much of the value of their work is independent of the particular language chosen.

Visits: OKI-Denki at Minato-ku, Tokyo

We visited an OKI lab in Minato-ku (Tamachi) which is involved with the ICOT program, on 30 October. Most of the day was spent in giving lectures on UW work in Fifth Generation computing; however we were shown OKI's SIM-P and given a short explanation. (We speculate that OKI lost the competition to build SIM-P prototypes for ICOT to Mitsubishi, but elected to go ahead with the project anyway. If true, this decision speaks volumes about OKI's opinion of the commercial potential of the ICOT program.

OKI's SIM-P was very much like Mitsubishi's, down to the same design limitation (some would say, design error). The machine was built of conventional SSI and MSI chips (except for 256K RAMs), in a lab which we considered to be poorly equipped. They have measured 10 000 LIPs (Logical Inferences/ second - a measure of speed for knowledge processing) and expect to get 20 000 - 30 000 LIPs. We felt that the OKI designers had not thought deeply or critically about the purposes of the SIM-P and hence about its design, but had just gone off dutifully to build hardware to ICOT's spec. (Meanwhile, of course, no-one at ICOT is completely sure that they've got the design right either ... ...)

OKI is also building a workstation with hardware support for object-oriented programming - a kind of successor to the interesting architecture but thoroughly botched implementation of the Intel 4/32. This work is almost certainly guided by Profs. Aiso and Tokoro of Keio University.

On 5 November, we returned to OKI to hear about their work to build a large, powerful Parallel Inference Machine (PIM) - the knowledge equivalent of a big mainframe. It will be a dataflow machine; the design was begun (at ICOT??) in April, 1984; they expect to bring it up in April 1985. Its behaviour has been simulated at functional level and its primitive operations have been simulated at cycle level. It will operate as a back-end processor attached to a DEC Vax, and will have both AND and OR - parallelism. Some of us felt that the design was ambitious, and that the project showed both boldness and courage. (The word used by another of us was "foolhardy".)

Presumably in response to the ICOT dictum, that the results of the project will (must?) be used to sharpen the tools available to the researchers - especially VLSI CAD tools - OKI has undertaken to do simulation of logic circuits (digital hardware, not to be confused with logic programming), using an object-oriented package written in Prolog. One of us being an old student of logic simulators, we can say with some certainty that the idea is silly, rather like cracking walnuts with a sledgehammer, and that the performance will be appalling unless they can somehow exploit the power of a dataflow architecture. Their presentation had an air of "We were expected to do something, and this is all we could think of" about it. The option of doing nothing was evidently unavailable.

OKI's commitment to the ICOT program seemed clear; they claimed to have four or five of "our best young engineers, unlike the third-raters which the other companies send", at ICOT. (Mitsubishi, interestingly enough, made the same claim.)

Fianlly, it was interesting to meet Computer Science postgraduate students from two of the better US schools (CMU and Princeton) working at OKI. Evidently they are part of a rotational or co-op scheme of some kind.

Visits: OKI-Denki at Takasaki

On 1 November we visited OKI's plant at Takasaki. The plant produces low-speed printers for personal computers (including IBM's PC) and is highly robotized. It produces a printer every 30 seconds using about 20 production employees. Two of these walk the floor looking for sick robots; the other eighteen do the final stage of printer assembly and test by hand. Without robots, 200 production workers would have been needed; the rather crude, non-sensory robots cost a few million dollars and paid for themselves in reduced labor costs in **about a year**. Better quality of product, tighter tolerances, and closer adherences to standards were the other benefits realized through robots.

There were also about one hundred employees doing **production engineering**; these jobs looked like a lot more fun than assembling printers must have been before robots. All of the engineering staff had engineering degrees or the equivalent of CEGEP diplomas.

During the visit we caught a glimpse of Japanese-style social relationships in the workplace. The Works Manager wore a cotton jacket furnished by the Company rather than an Italian-cut suit; he did not receive us in a fancy office, but in a rather austere meeting room, furnished about as lavishly as a university classroom.

(The sole exception was a **smart blackboard** which could produce an 8 by 11 paper copy of the material written on it, at the push of a

button. It could also be wired to a data network, for remotely-produced copies. The blackboard is an OKI product; we hope to get the first one sent to Canada.)

Visits: Fujitsu Ltd. at Kawasaki-City

We visited the Fujitsu Laboratories, directed by Hiroshi Yamada, at Kawasaki. We heard about their device work (the High Electron-Mobility Transistor or HEMT), supercomputer and gallium arsenide work. We also saw some natural-language work and saw a robot write Kanji characters on a single grain of rice. They seemed very impressed by the range and quality of work in our Logic Programming and Artificial Intelligence Group (the work is vertically integrated, from VLSI to Natural Language, similar to the work programs of some Japanese industrial groups). We detected less than overwhelming enthusiasm for the ICOT program. Fujitsu is often called "the IBM of Japan", so perhaps this attitude is understandable in those terms.

The Conference

The Conference consisted of two parts: progress reports by ICOT researchers, followed by research papers by workers from Japan and abroad. The progress reports confirmed, in detail, the conclusion reported here: Subject to minor deviations and re-definitions of objectives, the ICOT Program has achieved the objectives set for its First Phase (prototype and tool-building). However, many observers feel that the Second or Intermediate Phase (subsystems with problem solving, inference and Knowledge-base management functions) will be much harder to achieve, and will provide a more accurate indication as to the ultimate level of achievement of the entire ICOT Program.

The second part of the conference contained papers which exhibited the usual mix of topics and the usual variation in quality and relevance.

One of the few papers to report experimental results studied the speedup which might be obtained through AND, OR and unification parallelism. (These are fundamental operations of logic programming; OR-parallelism, for example, refers to the possibility of evaluating the terms of a logical disjunction (OR-operation) simultaneously (in parallel). The question is important, because this "evaluation in parallel" of the sub-operations within ANDing, ORing and unifying is a major strategy proposed for getting fast execution of logic programs and hence of knowledge-based computing.

The experiments reported were able to get at most 4 to 6 actions going simultaneously (i.e. were able to exploit a maximum of 4-6 central processing units), which yielded a speedup of 2 to 4 times, compared to the strictly sequential, one-operation-at-a-time mode of proceeding. This is hardly encouraging; in order to build large, powerful Parallel Inference Machines which can execute massive knowledge-based programs quickly, one would like to be able to exploit hundreds of central processing units. (Note that the success of the small, inexpensive Sequential Inference Machine - Personal or SIM-P does not depend on this question - the SIM-P ignores it.)

However, it is far from obvious that the objective of a large, high-speed, Parallel Inference Machine is essential to the larger goals of the Program. Large "mainframe" computers for conventional information processing are becoming less important due to the advent of desktop computers and networks of desktop computers; they same trend may occur in knowledge-based computing too.
Finally, the data are very preliminary. They may mean only that it is as easy for careless Prolog programmers to throw away the possibility of doing many things at once through thoughtless design, as it is for programmers who use other programming languages. In fact, this anecdote should be taken only as evidence to support the observation made above: The Second Phase of the ICOT Program, which includes the construction of PIMs, will be considerably more challenging than the First.

Conclusions and Recommendations

Our **Conclusions** are given in the Executive Summary at the beginning of this Report. Our **Recommendations** are that:

- **A Canadian program in this area of research is essential and must begin immediately.**

  Just as the UK, the USA, the EEC and the FRG have recognized the feasibility of Japan's goals, and the impact which the achievement of those goals would have on every area of human activity, so must we. The alternative is to slip slowly but steadily into the rôle of an under-developed nation.

- **The Canadian Program must ensure close co-operation between our Universities - our major source of expertise in this field - and private industry.**

  The universities are our major source of expertise in the technologies which underpin the Fifth Generation - microelectronics, computer architecture, software engineering and artificial intelligence. They are also our only source of the trained people who are and will be in critically short supply. Without the involvement of private industry to create wealth-producing, socially beneficial new products and services, the whole exercise would be pointless.

• **The Natural Sciences and Engineering Research Council (NSERC) should be asked to organize the Program.**

NSERC is the major sponsor of university research in natural sciences and engineering, which includes all of the topics relevant here. NSERC's use of peer review and its administrative structure have the full trust and confidence of Canadian academics, and its Industry-University programs have been instrumental in bringing about closer cooperation between university researchers and Canadian industry. Finally, NSERC has sponsored the creation of the Canadian Society for Fifth Generation Research, to advise the Council on questions related to Fifth Generation computing research.

Appendix:
Materials Collected

Copies of these materials, which include

- glossy handouts from the companies visited
- visual aids used by researchers at OKI-Denki and Mitsubishi
- a listing of system calls supported by the Mitsubishi SIM-P operating system
- material describing the ICOT program

are available on request to

Institute for Computer Research
MC6018
University of Waterloo
University Avenue
Waterloo Ontario
N2 L3 61

telephone: 519 885 1211 X3509
Usenet: watmath!watdaisy!egmanning