

Some Reflections on the Fifth Generation Project

Seif Haridi & Siwert Sundström
Swedish Institute of Computer Science

This report represents the personal opinions of the authors.

Summary

The Fifth generation project has been a very valuable project. ICOT has served as a source of inspiration to international research in the area. It is not unfair to say that many research programmes and institutes were conceived in reaction to the establishment of ICOT. ICOT has made a solid progress in applying parallel processing to symbolic computing and knowledge information processing. Many of the ideas and concepts developed by ICOT have been further developed and modified at other research centers. We believe that ICOT has contributed to and affected other research more than outside research has affected and influenced ICOT. This may have been inevitable due to the fixed nature of ICOT's long-term plan. Therefore, we must also look outside ICOT to judge the results of the project.

The internal technical results, given the constraint of a fixed time period, and the temporary nature of its personnel, surpass expectations. However, we find that at this stage the project remains unfinished.

In order to fully achieve the goals of the project, research should continue on improvement of the software generated by the project and on careful evaluation of the software and the parallel architecture prototypes. It is also of extreme importance to port the valuable software generated by the project to widely accessible general purpose machines, both parallel and sequential. This is necessary in order to disseminate the results of the project to the international research community and to create a basis for possible commercial exploitation. Experience from developing knowledge processing tools and databases has to be documented and fed back to improve the design of KL1. In particular, it seems that there is a need to improve the process communication abilities of KL1. There is also a need to incorporate recent research results on integrating the programming paradigms of Prolog into the language as shown in the Andorra family of languages. This is vital to facilitate the development of knowledge processing tools. Finally there is a need for research on software tools that integrate

both knowledge information processing and numeric computations.

The Goals of the Project

It is the understanding of the authors that the ICOT project was set up with the following combined goals:

1. To develop the knowledge processing technology and the basic tools needed for advanced applications that require both deductive and intensive symbolic processing capabilities. These applications were predicted to be increasingly in demand during the nineties and early 21st Century.
2. To exploit and develop parallel processing technology and highly parallel architectures in order to meet the performance requirements expected for the above class of applications.
3. As the link or glue between the higher level knowledge-based applications and the highly parallel machine at the lower level, an intermediate programming language based on logic programming was to be developed. This language, on one hand, would be flexible enough to implement the high level tools needed for the knowledge intensive applications, as stated in point 1, and, on the other hand, be suitable for exploiting the resources of parallel machines and making parallel programming easy as stated in point 2.

Needless to say the goals of the project at the time of its conception were very ambitious, since they required breakthroughs in many of the technological areas outlined above. However, that is not to say that they were not viable. In fact, with the current knowledge of the state of art, these goals seem achievable, and we certainly admire the courage of the project leaders to attack such challenging technological problems at such an early stage.

Organisational Structure Tradeoffs

The project leadership was given ten years to achieve these goals, and the majority of the staff were recruited from the Japanese computer industry on a temporary basis. Under these constraints, the ICOT project can be evaluated from two different perspectives, one w.r.t. its scientific output, and the other w.r.t. its industrial output. It is, however, important to notice that under the premises of the project, i.e fixed time and temporary personnel, the goals of maximising both the scientific and industrial outputs are largely incompatible and some compromises had to be made.

Just to illustrate some of the trade-offs, here are some issues:

1- To maximise the scientific output of the project, it is necessary to have well trained key persons stay on after spending their assigned period in the project (approximately 4 years) while, in contrast, to maximise the industrial benefits of the associated companies technology transfer, through experienced personnel transfer, should occur regularly. The policy used in the project, as far as the authors are aware, was something in between; most of the key persons stayed in the project while some left the project after their assigned period.

2- To fix a project period is extremely counter-creative and correspondingly counter productive as regards to the scientific output of the project. It is better in scientific endeavours to organise projects into stages, where the financing of the next stage is conditional on the progress made in the previous one, and to set the goals of the next period depending on the progress of the previous one, and possible changes in the environment of the project. This would make projects more flexible, and responsive to unexpected changes within the project and its surroundings. Moreover, if a project, under a fixed time constraint, has some concrete technological goals to achieve, a sizeable portion of its late period will inevitably be spent in pure development efforts, thus missing some opportunities of fundamental advances in the state of art during this late period. It is again the view of the authors, that the ICOT project, during its initial and intermediate period, played an important role in advancing the state of art towards achieving the goals of the project, while the last 4 years were mostly spent in developing the prototypes of the technology envisioned at the intermediate phase.

3- A scientific research project would benefit from open cooperation with other researchers and research organisations, while an industrial project is normally closed for competitiveness reasons. The ICOT project was more a scientific project than an industrial one in this respect, but not entirely. One important obstacle in this case is that most of the software developed by the project was on proprietary hardware which effectively prevented other researchers from sharing the results of the project, and resulted in the missing of opportunities that would otherwise have occurred through early feedback and valuable evaluation reports from other researchers. ICOT tried to remedy this problem by distributing some of its hardware to other research organisations but such hardware was not in fact effectively used for different reasons. Exchange of software among researchers is a very important factor in improving the quality of the research produced by any computer science research organisation. MITI lately released the software of the project in the public domain, but this should have happened earlier.

Scientific achievements of ICOT

On the positive side

Initially at the start of the project there was a clear world-wide trend to use dedicated machines for AI, and knowledge intensive applications. ICOT responded successfully

to this trend by developing the PSI machine, and most importantly by developing ESP. ESP was the first programming language that combined both the logic programming paradigm as embodied in Prolog and the object oriented programming paradigm. The flexibility of the language was clearly demonstrated by programming the operating system of PSI and developing a number of tools for knowledge representation and natural language processing.

ICOT played a central role in the development of programming languages for parallel symbolic computing. The design of GHC was certainly a very important step in the development of concurrent logic programming languages.

ICOT has contributed considerably to the implementation technology of programming languages on parallel architectures. Techniques for memory management like the MRB scheme and parallel garbage collection on cluster basis, were certainly innovative techniques developed by ICOT.

ICOT has also contributed considerably to the development of constraint logic programming by for example, the work on CIL and CAL.

The development of PIMOS is the first large scale operating system for parallel computers.

ICOT demonstrated, beyond doubt, the usefulness of parallel processing in symbolic and highly irregular computations.

On the negative side

ICOT invested a lot of effort trying to compile the programming paradigms of PROLOG to KL1, but one must say that these efforts have only partially succeeded. Research done by outside researchers has shown that it is possible to combine Prolog and GHC (a good example is the work done at SICS on Andorra). However, the results came too late to be exploited within the remaining time allocated to the project. This is very unfortunate since a lot of ICOT's early research on natural language processing and constraints was based on Prolog and could not be exploited effectively when the project moved from ESP to KL1. Moreover there is some doubt about the flexibility of the communication structures of KL1.

Due to the lack of time the results of ICOT on knowledge representation and knowledge bases remain unevaluated by the outside research community.

It is questionable whether the architectures produced by the project are adequate for efficiently implementing languages like KL1. The most important issue is the lack of support of global address spaces, and multiple-users address spaces in the hardware. It is widely accepted by researchers in computer architecture that global address space

is a very important factor in reducing the problems associated with dynamic data partitioning and dynamic load distribution.

Recommendation

Most of the problems mentioned above can be remedied if ICOT's research activities will continue. Therefore we recommend that ICOT:

- should perform research on improving KL1 to integrate don't care and encapsulated don't know nondeterminism.
- should carefully evaluate and improve the knowledge representation tools through feedback from sizable applications.
- should evaluate the parallel PIM machines to identify the minimal hardware requirements needed to efficiently execute parallel symbolic languages like KL1.
- should investigate the possibility to integrate more flexible communication facilities in KL1 (e.g. ports in the Andorra Kernel language), and/or the ability to handle distributed arrays efficiently (e.g. M-structures in dataflow languages).

Industrial achievements of ICOT

On the positive side

It is clear that a large number of young engineers and scientists from the associated companies have been trained in advanced technology, which would not have been possible without the ICOT project; in particular in areas of, (i) design, development and implementation of high-level programming languages,

1. knowledge representation languages, (iii) scalable operating systems and
2. design and implementation of multiprocessor architectures.

The knowledge and the skills developed while working on the cache coherent clusters of PIM/p, PIM/c and the two level cache coherent system of PIM/k must have been very important since the current commercial trend of successful parallel machines is cache-coherent ones. We congratulate the project leadership on producing the hardware prototypes on that scale given the tight time constraints.

It is a great achievement to be able to deliver on time, nearly all what has been promised when the project started 10 years ago, together with impressive number of demo applications running on parallel machines.

On the negative side

Since 1985 there was a clear shift from dedicated hardware to general purpose hardware. ICOT was not flexible enough, perhaps due to the overall project structure. A very valuable possible industrial output would have been setting standards for the design of general purpose RISC chip for symbolic languages. This might have been just adding tag support for conventional RISC chips, and/or support for multiprocessing. Another issue is the lack of active and incremental migration of ICOT's software to commercially available machines in general, and to machines belonging to the associated companies in particular. This is crucial for any commercial exploitation, and is a prerequisite for more eager adoption of new technology within the companies engaged in the project.

Recommendations

ICOT software should be ported to a wide class of commercially available machines in order to allow for adoption of the technology produced within the associated companies and the wide research communities.

Minimal extensions to existing standard hardware to efficiently support parallel symbolic computing should be identified, and reported to the associated companies to be taken into account when new generation of commercial hardware is designed by the associated companies.

Conclusions

The Fifth generation project has indeed been a very valuable project. From a Swedish perspective ICOT has succeeded in establishing close cooperation between Swedish and Japanese researchers both on the academic and industrial level. ICOT has also influenced the establishment of the Swedish Institute of Computer Science.

The internal technical results, given the constraint of a fixed time period, and the temporary nature of its personnel, surpass expectation. However, it is clear that at this stage the project remains unfinished.

In order to fully achieve the goals of the project, research should continue on improvement of the software generated by the project and on careful evaluation of the parallel architecture prototypes. It is also of extreme importance to port the valuable software generated by the project to widely accessible general purpose machines, both parallel and sequential. This is necessary in order to disseminate the results of the project in the international research community, to stimulate the adoption of this new technology within the associated companies, and to create a basis for possible commercial exploitation. Experience from developing knowledge processing tools and

databases has to be documented and fed back to improve the design of KL1. In particular, it seems that there is a need to improve the process communication abilities of KL1. There is also a need to incorporate recent research results on integrating the programming paradigms of Prolog into the language. This is vital to facilitate the development of knowledge processing tools.

Siwert Sundström, born 1933 in Sweden.

Brief resume

- 1960 Royal Institute of Technology, Stockholm, Sweden
Masters degree in Electrical and Electronic Engineering
- 1960-1964 Facit Electronics Co., Solna, Sweden.
Computer hardware design.
- 1964-1965 North Electric Co., Ohio, USA
System studies in the field of Electronic Telephone Switching Systems
- 1966-1978 Ericsson Telecommunications Co., Stockholm, Sweden
Starting with system development of Control system for Electronic Telephone Switching Systems
Ending as department manager responsible for Ericssons first generation computer controlled telephone switching system.
- 1978-1985 Ellemtel Development Co., Stockholm, Sweden
Starting as department manager, responsible for development of the Digital Telephone Switching System, AXE.
Ending as department manager, responsible for Software Design Support Systems.
- 1985- Swedish Institute of Computer Science, Stockholm, Sweden
Managing Director
The institute is a non-profit research organization jointly financed by industry and government. 60 permanently employed researchers.
Research fields:
 Distributed Systems,
 Logic programming and Parallel computer Systems
 Knowledge Based Systems.