General Report of the FGCS Follow-on Project

Shunichi Uchida

Institute for New Generation Computer Technology
4-28, Mita 1-chome, Minato-ku, Tokyo 108, Japan
uchida@icot.or.jp

Abstract

This paper is to report activities of the two year FGCS Follow-on Project which was carried out after the eleven year FGCS Project.

The activities in the Follow-on Project can be roughly divided into two kinds. One is research and development and the other is dissemination of the FGCS technology.

The research and development activities include the development of a new KL1 programming environment called KLIC on Unix-based sequential and parallel machines, and the modification and extension of major knowledge processing tools and application systems to port them to these machines using KLIC.

The dissemination activities include organization of various collaborative research projects with domestic and foreign research institutes and universities, holding workshops and seminars, distribution of software products as ICOT Free Software, and technical information on research and development activities by the anonymous FTP and World Wide Web (WWW) on the Internet.

Now, KLIC is operational on dozens of Unix-based workstations and several MIMD parallel machines available on the market. Several parallel tools and application systems have already been ported to these machines and used with various existing software systems available in Unix environments.

All of these tools and application systems are distributed as ICOT Free Software. Since its first release in August 1992, more than 2,100 people have transferred more than 12,000 files. KLIC has already been used for teaching parallel languages in dozens of universities in Japan and overseas.

We are now confident that the Follow-on Project will attain its general goal and successfully finish in March 1995. After the ICOT research center closes in March 1995, we plan to continue maintenance and distribution of ICOT Free Software for a few years.

1 Introduction

Recently, the main stream of computer science and technology is rapidly changing toward parallel and distributed computing. This change could be regarded as the opening of a new era of computers. It started from hardware technology and is now spreading to software technology and applications.

In this new era, FGCS technology is becoming a more important and practical key technology which provides us with the capability of organizing larger-scale parallel and distributed knowledge processing systems.

The successes of the FGCS Follow-on Project has expanded the use of the FGCS technology and will augment its contribution to the world-wide advance of computer science and technology.

This paper reports the activities of the two year FGCS Follow-on Project which was started in April 1993 after the eleven year FGCS Project.

The Follow-on Project aims to disseminate the FGCS technology which was developed in the FGCS Project by combining highly parallel processing and knowledge processing technologies using logic programming as its kernel language.

The central portion of the FGCS technology is the KL1 parallel language and PIMOS operating system implemented on parallel inference machines (PIMs) having about 1000 element processors.

Using this KL1 and PIMOS environment on PIMs, many parallel software systems have been developed. Many of them, such as theorem prover MGTP, have proved that the KL1 and PIMOS environment has enabled us to gain a linear speed-up which is almost proportional to the number of element processors.

This is probably the first time in the world that major knowledge processing or AI application systems could gain an almost linear speed-up using highly parallel processing. The KL1 and PIMOS environment running on the PIMs is now the most powerful parallel system for knowledge processing.

At the end of the FGCS Project, ICOT and the Ministry of International Trade and Industry (MITI) thought that the FGCS technology should be disseminated as a

Table 1: Change of Budget and Number of People from fiscal '91 to '94

Project name	*Fiscal year Apr.to Mar.*	*Budget billion yen*	*ICOT members*	*Subcontracted company members*
FGCS (10th year)	'91	7.2	90	500
FGCS (11th year)	'92	3.6	60 ('92/10)	200
Follow-on (1st year)	'93	1.3	40 ('93/ 4)	50
Follow-on (2nd year)	'94	1.3	40	50

new common infrastructure for advanced research into computer science and technology. For this reason, ICOT placed all the major software developed in the FGCS project in the public domain as ICOT Free Software.

However, all the major software systems including the PIMOS operating system are written in KL1 and able to run only on PIMs. This is obviously an obstacle to the dissemination.

Around the end of the FGCS Project, general purpose MIMD parallel machines started to appear in the market for large-scale numerical processing applications. The scale of these parallel machines is not as large as the PIMs. However, they have the potential to provide us with more processing power at low cost in the near future. Furthermore, they are equipped with a Unix-based operating system with some extension for parallel processing such as the software called "Parallel Virtual Machine" (PVM).

Considering this situation, ICOT and MITI decided to carry out the FGCS Follow-on Project. The primary technical goal of the Follow-on Project is to develop a new KL1 and PIMOS environment on Unix-based parallel machines. This new environment is named KLIC because a KL1 program is compiled into a C program in this environment.

The secondary goal is to port several new and interesting systems developed in the FGCS Project to Unix-based machines using KLIC. To achieve this goal, these systems had to be made compact so that they could run on the Unix-based machines which are currently much smaller than the PIMs. Furthermore, new user interface portions have to be developed using standard software tools such as X-windows Motif.

These goals have almost been attained now. Some of new software systems have already been distributed as ICOT Free Software. The rest will be completed and started to be distributed in March 1995 by the anonymous FTP and World Wide Web (WWW) on Internet.

These software systems including the KLIC are now used not only for practical applications but also for education of parallel languages, programming methodology and parallel knowledge processing in universities. We are sure that the more inexpensive and popular large-scale parallel machines become, the greater the contribution of the FGCS technology will be.

2 Research and development activities

2.1 Starting the Follow-on Project

In June 1992, the decision to carry out the Follow-on Project was made and announced at the FGCS'92 international conference. We then started the preparation for the Follow-on Project.

We had determined an outline of the technical goals for the Follow-on Project, but detailed research topics had to be decided based on whether we would be able to keep the relevant researchers at the ICOT research center.

In the final stage of the FGCS Project, many of the ICOT researchers had been at ICOT for 4 or 5 years. This length of stay was too long compared to the standard period for rotating researchers between ICOT and their companies. Thus, we had to let many of them return to their parent companies and ask the companies to send new younger researchers to ICOT. We negotiated with both the ICOT researchers and their companies.

As the idea of a Follow-on Project was well received by researchers and companies, we were able to keep research leaders in the important research topics for the Follow-on Project and recruit about 20 new researchers. Many of the new researchers had been involved in the FGCS Project in their companies or software houses and had sufficient knowledge and skill to perform their assigned job.

Table 1. shows the changes in budget and staffing levels from fiscal '91 to '94.

In fiscal '92, the number of ICOT researchers decreased from 90 to 40 and the 7 research sections were reduced to 2 research departments. In affiliated companies and software houses, the number of researchers, engineers and managers decreased from 500 to 50. These 50 people were brought to work at the ICOT research center.

While many researchers and engineers left the FGCS Project, about 30 researchers who had been involved in the FGCS Project at ICOT or in their companies moved to major universities between fiscal '91 and '93. Many of them received Ph.D's for their research work with the FGCS Project. These researchers are playing a very important role in the Follow-on Project.

2.2 Selection of research themes

As the budget and the number of ICOT researchers shrank, we had to curtail research themes. In the final stage of the FGCS Project, we had more than 50 research themes, including about 20 parallel application systems. These included practical research themes such as expert systems for job scheduling and VLSI CAD systems, as well as basic research.

The general technical goal of the Follow-on Project is to make major software systems developed in the FGCS Project operational on Unix-based parallel machines. Thus, the primary technical goal is the development of a new KL1 and PIMOS environment named KLIC on Unix-based machines.

In choosing other themes, selection criteria included whether the theme would have a large impact on the future of computer science, and whether progress could be effectively accelerated by the use of parallel processing.

We chose the following research themes for the Follow-on Project. They are divided into two groups.

- Parallel Basic Software
  1. KLIC system: a KL1 programming environment for sequential and parallel Unix-based machines
  2. Evaluation of PIM architectures and their KL1 language processors
  3. Parallel nested relational DBMS, Kappa
- Knowledge Processing Software
  1. Parallel theorem prover, MGTP
  2. Knowledge representation languages;
     - Deductive and object-oriented language, $\mathcal{QUIXOTE}$
     - Parallel constraint logic programming language, GDCC
     - Heterogeneous distributed problem solving system, Helios
  3. Genetic information processing systems
     - DNA and protein sequence alignment and editing system
     - New algorithms for sequence and structure analysis
     - Biological DBMS and KBMS
  4. Legal reasoning system, new HELIC-II

2.3 Research activity of each theme

2.3.1 KLIC system

The KLIC system consists of a KL1 compiler and a run-time library. The KL1 compiler is written in KL1 and compiles a KL1 program into a C program. The run-time library is prepared as a library of C programs which provide functions such as debugging, monitoring, parallel execution management, resource management and so on. These functions are almost the same as the ones which PIMOS provides on the PIMs.

Some of these C programs are linked to a user program, if it uses some of these functions. If the user program does not use any of these functions, it is compiled into a very simple C program which is executable even on a small personal computer. This nature of KLIC is very convenient for educational purposes.

The development of KLIC was done in two steps. In the first step, a KLIC sequential version was developed. Development of the KL1 compiler was a main topic. This version could attain a fairly good performance. It attained 2 MLIPS on an SS-10/30 and 3.7 MLIPS on DEC AXP. The KLIC sequential version was released in November 1993 as ICOT Free Software.

In the second step, a KLIC parallel version was developed. This version is more complicated than the sequential version. It uses the software called PVM which provide us with a standard interface for inter-processor communications over parallel or distributed machines. This software is distributed as PDS and is provided for most of the recent parallel machines.

A new KLIC run-time library for the parallel version was developed to provide functions for parallel execution management, resource management, debugging and monitoring, and so on. First release of the KLIC parallel version was made in September 1994 and will be released in February 1995 as ICOT Free Software.

The KLIC parallel version will be ported to parallel machines such as Sparc Center, DEC AXP 7000, CM-5, AP-1000+, Cenju-3, SR-2001, SP-2 and so on.

2.3.2 Evaluation of PIM architectures

This theme intended to further investigate five models of PIMs. All models support KL1; however, the architecture of the processing element and inter-processor connection mechanism in each model are different.

Designers of the five PIM models formed a task group and had interesting discussions. They gathered evaluation data and compared unique design features with each other. The general evaluation results are presented at the FGCS'94 symposium.

2.3.3 Parallel nested relational DBMS, Kappa

Kappa is an extended relational DBMS which permits us to use nested tables to represent data.

In a standard relational DBMS, we usually use two dimensional tables to represent data structures. One table is divided into many rectangular boxes by rows and columns. The structure of this table is very regular and rigid. We are permitted to put only one data item in each box. This regularity is an advantage as long as we use only simple data structures.

A nested relational DBMS permits us to put multiple data items in one box. Thus, it is beneficial to handle complex data structures such as natural language dictionaries, biological data, and so on. We can represent complex data in more compact and comprehensive format in a nested relational DBMS than in a standard relational DBMS.

We developed Kappa on the PIM and made full use of parallel processing to gain better performance.

In the Follow-on Project, Kappa was almost reconstructed to be a compact and faster system. Some of its low-level functions were re-written in C code for better performance.

Now it can attain an almost comparable performance to ordinary relational DBMSs for a complex database such as GenBank on a sequential machine. If it is used on a parallel machine, it will attain even better performance using parallel processing.

2.3.4 Parallel theorem prover, MGTP

MGTP is a model generation theorem prover for full first-order logic. It is one of the most successful application programs in the FGCS Project. It demonstrated the effectiveness of the KL1 and PIMOS environment on the PIMs.

Generally, it has been well known that theorem provers have a very large search space and thus would be an interesting application of parallel processing. However, its computational structure is very irregular and it is hard to predict how its search tree will extend its branches.

Thus, we need to divide the computation into parallel-processable processes and their allocate them to many element processors dynamically on the fly. It is very hard job to do this with conventional methods.

MGTP successfully implemented this job division and allocation in the KL1 and PIMOS environment on the PIM model in late 1991. The execution speed increased almost proportional to the number of element processors.

Furthermore, the program production period was surprisingly short. This was the first clear evidence that the KL1 and PIMOS environment on PIMs can successfully be applied for solving a large-scale AI problem.

The experiences and skills of the MGTP group were quickly transferred to other research groups at ICOT. This encouraged all the researchers at ICOT and their companies.

In the Follow-on Project, the MGTP group further developed their provers. As a result, MGTP became the fastest theorem prover of this kind in the world and proved some open problems in quasi group theory.

Furthermore, some of the MGTP provers were provided as tools for practical applications such as the rule-based engine of the HELIC-II legal reasoning system.

This indicates that a theorem prover can be regarded as a higher-level inference engine and that it can be adapted for knowledge processing applications such as KBMS, natural language understanding, and software engineering.

In addition to the ordinary MGTP provers, the MGTP group for educational purposes is developing compact versions of their provers using Prolog.

2.3.5 Knowledge representation languages

In the FGCS Project, the research on knowledge representation languages at ICOT has been based on mathematical logic. The research started from nested relational databases and deductive languages. In the intermediate stage, research on constrained logic programming languages was started.

On the other hand, research on object-oriented languages (O-O languages) at ICOT was mainly done with system description languages such as ESP, which is extended Prolog developed as a programming language for PSI machines.

Outside of ICOT, some database researchers were attracted by O-O languages and their mechanisms for modularization and inheritance. They developed O-O databases and indicated that they are able to provide us with more flexible data models than relational databases and deductive databases. We started discussions on how to introduce the merits of O-O databases into our framework.

In the final stage of the FGCS Project, we started the design of a deductive and O-O database language and DBMS. Finally, the design of this language and system was completed and became the knowledge representation language, $\mathcal{QUIXOTE}$. The first version was implemented in KL1.

As $\mathcal{QUIXOTE}$ has rich O-O based functions combined with a deductive language, it can fulfill requirements for describing complex knowledge fragments such as legal rules and biological reactions. However, because of the wealth of its functions, implementation was so complicated that performance and the stability of the software system were not satisfactory.

In the Follow-on Project, intensive efforts were made to improve language the specification and system implementation, and the system was ported to Unix-based machines using KLIC. $\mathcal{QUIXOTE}$ will be released as a more practical system in March 1995.

This system is now called the "big"-$\mathcal{QUIXOTE}$. Micro-$\mathcal{QUIXOTE}$, which is a compact subset was also developed for educational purposes. Micro-$\mathcal{QUIXOTE}$ is small enough to run on a personal computer.

2.3.6 Genetic information processing systems

Research on genetic information processing was started with the parallel processing of multiple alignment of protein sequences. This research topic is continued in the Follow-on Project and has now been extended to a sequence alignment and editing system.

This system can handle both protein and DNA sequences. Its alignment algorithm is based on a DP matching algorithm and implemented in KL1. Parallel implementation of this algorithm has been improved many times to attain better performance on a PIM.

Recently, the use of a genetic algorithm has given us better alignments for some interesting cases. The use of constraints between some amino acids or nucleic acids was tried in order to narrow the search space. This system is now being ported to Unix-based parallel machines and is available to biologists.

Research on the prediction of protein structures has been made in the Follow-on project. The use of the Hidden Markov Model (HMM) has given us interesting results in some cases.

Research on biological DBMS and KBMS has also been made in connection with research on knowledge representation languages. Some biological reactions as well as the characteristics of biological and chemical materials are written in *Quixote* and stored in the knowledge base.

Reactions were coded into hundreds of figures by biologists. They are now described in *Quixote* and can be retrieved by the KBMS. We expect that this kind of use of knowledge representation languages will be beneficial in many other application areas.

In the research on genetic information processing, research collaboration is essential. We have set up several international and national collaborative research projects with biologists and computer specialists. They have brought to us many fascinating problems in recent molecular biology. We hope that the software tools we produced will continue to assist them with their research.

2.3.7 Legal reasoning system, new Helic-II

Research on the Helic-II legal reasoning system was started in the final stage of the FGCS Project. It is one of the application systems developed to generally evaluate the KL1 and PIMOS environment, knowledge representation languages and other software tools.

This sytem was very successful not only in demonstrating the usefullness of the FGCS technology, but also in showing us possible uses for the FGCS technology which are beyond our initial expectations. This system gives us a better understanding of how FGCS technology can be used for applications in social-scientific areas.

To analyse a given case and predict all the possible judgements, Helic-II used two knowledge bases and two inference engines. One is a case knowledge base combined with a case-based inference engine. The other is a rule knowledge base combined with a rule-based engine. The rule-based engine was built using the MGTP theorem prover as it kernel.

To describe the penal code and case rules, some of the results of the research into knowledge representation languges were used. Both inference engines were implemented in KL1 and were speeded up by parallel processing on a PIM.

As Helic-II used many other research results developed in the FGCS Project effectively, it was the most appropriate program for general evaluation of the FGCS technology.

In the Follow-on Project, this research is progressing to include more sophisticated functions such as simulating a debate between a prosecutor and lawyer. This extended version is called the new Helic-II system and requires descriptions of more cases and rules.

To do this, more background knowledge on the law and legal systems is vital. Thus, collaboration with researchers and experts in a low area is indespensable. Setting up this collaboration needed more effort than for genetic information processing because of the wider culture gap between computer science and law.

The ultimate goal of this research is still very far away, but the spin-offs from reseach are expected to be used in many social scientific applications where the problems to be solved will be simpler than this system. The new Helic-II is being ported to Unix-based machines and released as ICOT Free Software.

3 Activities for dissemination

3.1 Framework for dissemination

Around the end of the FGCS Project, MITI organized a high-level committee to assess the research results of the FGCS Project. [DPBCT 1994] One of its conclusions was that the results were considered to be still so far away from the market's needs that computer companies could not commercialize them in a few years although they can be highly evaluated from the academic viewpoints.

Thus, dissemination activities were mainly directed at academia, that is, national and foreign research institutes and universities. However, quite recently, computer companies both of home and abroad have become more positive and supportive of our dissemination activities.

We prepared for dissemination in two aspects. For the first one, we prepared hardware and facilities such as file servers connected to the Internet accompanied by a team to maintain the programs and documentation of ICOT

Free Software, and high-speed network links so that our research partners can use the PIMs from remote sites.

For the second one, we have put our efforts into networking among ICOT members. We have been able to keep some of the collaborative research projects which were set up in the FGCS Project and have tried to set up new ones with national and foreign institutes and universities based on the new research themes of the Follow-on Project.

3.2 Collaboration on research in Japan

3.2.1 Organization of Task Groups

In the FGCS Project, we had several "Working Groups (WGs)" which recruited researchers from academia and industry and provided them with the opportunity to exchange technical information and provided some research tools at ICOT.

The main activity of the WGs was to meet and discuss specific research topics and so the relationship between the WG members and ICOT was not so tight.

In the Follow-on Project, we needed a tighter relationship with some researchers in universities and research institutes in order to to ask them to carry out parts of our research and development or to help us hold seminars and workshops and so on.

Thus, we set up new Task Groups (TGs) to recruit researchers from universities and industry. Naturally, many of them had been involved in the FGCS Project at ICOT or in their companies.

They are playing a very important role in the dissemination process. For example, we have held KLIC seminars several times at universities with the help of these TG members.

We currently have the following 7 task groups. The numbers shown in parentheses are the regular members excluding ICOT researchers.

1. Parallel Symbolic Processing System Task Group (18)
2. Parallel Inference Machine Evaluation Task Group (10)
3. KLIC Task Group (9)
4. Parallel Theorem Proving Task Group (9)
5. Heterogeneous Knowledge-Base Task Group (17)
6. Protein Structure Prediction Task Group (17)
7. Legal Reasoning System Task Group (7)

3.2.2 Collaboration with Japanese universities and research institutes

Some of the research results of the FGCS Project are practical but some others are not. KLIC is an example of a practical one. Results are integrated into software systems which can be used as useful tools. Thus, technology transfer is rather easy.

However, some others were not well integrated into complete software systems. For example, parallel programming methodology using KL1 is partially described in technical papers and partially integrated into some KL1 programs, but mostly existing in the brains of the researchers.

Results have not been well ordered yet and thus are very abstract. Experience on how to represent legal knowledge gained in the development of the legal reasoning system is another abstract result.

To disseminate these abstract research results, we decided to set up several small collaborative projects with Japanese universities. In these projects, professors are expected to use some of our results as seeds in starting new research projects. Of course, they are asked to use practical results as educational tools or as new infrastructures.

Currently, we have the following 15 projects.

1. Parallel language processors and environments;
   - Optimization of a language processor (Prof. H. Tanaka, Univ. of Tokyo)
   - Optimal implementation of KL1 (Prof. H. Nakashima, Kyoto Univ.)
   - Program analysis and optimization (Prof. K. Ueda, Waseda Univ.)
   - Visual interface of parallel systems (Prof. J. Tanaka, Tsukuba Univ.)
2. Parallel theorem proving;
   - High-speed inference (Prof. K. Fuchi, Univ. of Tokyo)
   - Lannguage processor based on theorem proving (Prof. M. Amamiya, Kyushu Univ.)
3. Natural language processing;
   - NL processing on the PIM (Prof. H. Tanaka, Tokyo Institutes of Technology)
   - Parallel NL understanding (Prof. R. Taniguchi, Kyushu Univ.)
   - NL tools (Prof. Y. Matsumoto, AIST-Nara)
4. Others;

- Distributed AI systems
  (Prof. K. Furukawa, Keio Univ.)
- Constraint processing
  (Prof. F. Mizoguchi, Science Univ. of Tokyo)
- Parallel VLSI CAD systems
  (Prof. K. Taki, Kobe Univ.)
- Evaluation of parallel architectures for LP
  (Prof. H. Tanaka, Univ. of Tokyo)
- Transaction management
  (Prof. H. Yokota, JAIST-Hokuriku)
- Methodology for parallel scientific computation
  (Prof. S. Kunifuji, JAIST-Hokuriku)

Furthermore, we have two collaborative projects with the Electrotechnical Laboratory (ETL) and the Mechanical Engineering Laboratory (MEL). These are national laboratories which belong to MITI. These projects are continuations of the FGCS Project.

3.3 Collaboration on international research and achievement sharing

The technical goal of the FGCS Project was very advanced and needed the collaboration of the world's leading researchers. Thus, international collaboration began in the initial stage of the FGCS Project. Early collaboration was based on invitation or mutual visits of individual researchers.

From the intermediate stage, we concluded formal agreements for exchanging researchers and holding workshops with the following governmental organizations;

- National Science Foundation (NSF) in the U.S.A.
  (Mr. Y.T. Chien)
- Institute National de Recherche en Informatique et en Automatique (INRIA) in France
  (Dr. L. Kott and Dr. G. Kahn)
- Swedish Institute for Computer Science (SICS)
  (Dr. S. Haridi and Dr. M. Nilsson)
- Department of Trade and Industry (DTI) in the UK
  (Dr. K. Shotton and Dr. P. Rothwell)

In the final stage, we started more substantial research collaborations to find appropriate applications for the KL1, PIMOS and PIMs. We exchanged ideas and research tools with the following U.S. laboratories;

- Argonne National Laboratory (ANL)
  (Dr. E. Lusk, Dr. R. Overbeek and Dr. R. Stevens)
- National Institute of Health (NIH)
  (Mr. R.J. Feldmann and Dr. G.S. Michaels)
- Lowrence Berkeley Laboratory (LBL)
  (Dr. C. Cantor and Dr. C. Smith)

With researchers at ANL, we conducted research on biological analysis and theorem proving. They gave us several interesting problems on theorem proving. They greatly stimulated our research on parallel theorem proving and biological analysis. We held US-Japan workshops on logic programming and theorem proving twice at ANL sponsored by NSF and ICOT.

With researchers at NIH and LBL, we conducted research on protein structure analysis and sequence alignment. We learned much about molecular biology through these collaborations.

The collaborative research on theorem proving was extended to include Australian researchers.

- The Australian National University (ANU)
  (Prof. M. McRobbie and Dr. J. Slaney)

ANU worked with us to try to solve some open problems in the quasi-group theory using the MGTP theorem prover on a PIM. This trial was very successful. The achievements of this collaboration stimulated the world's researchers in this area very much. We held several international workshops with these researchers to share our research results and experiences.

These collaborations were really helpful to our growing research groups into genetic information processing and parallel theorem proving. Most of these collaborations are still continuing on an individual basis.

In the Follow-on Project, we set up new tighter collaborative research projects with the following two universities;

- The University of Bristol
  (Prof. D. Warren and Prof. S. Gregory)
- The University of Oregon
  (Prof. J. Conery and Prof. E. Tick)

We hoped to further develop the KLIC system and application systems as well as the disseminating the KLIC system and other programs as ICOT Free Software.

With researchers at the University of Bristol, we are working to develop a parallel debugger and a constraint solver which can be used as extended functions in the KLIC system.

With researchers at the University of Oregon, we are conducting research on optimizing the compiler for KLIC and biological sequence analysis using constraints. Based on this collaboration, we held the fifth US-Japan workshop at this University attended by international researchers.

The results of all these collaborations will be placed in the public domain and distributed as extensions of ICOT Free Software.

With the completion of the Follow-on Project, we have to terminate all of these collaborations. We hope that the work will continue on individual bases.

3.4 Distribution of ICOT Free Software (IFS) and other technical information

We started the distribution of ICOT Free Software (IFS) using anonymous FTP on the Internet in August 1992. Since then, 12,000 files have been transferred to 2,200 sites in 45 countries. Among these file transfers, 40 % were to Japan, 30 % to the U.S.A. and the remaining 30 % to other countries.

In the final year of the FGCS Project, we placed 77 programs in the public domain as IFS and started up ICOT's FTP server in August 1992. These programs are relatively large. For example, the PIMOS operating system is counted as one program. These programs include many large programs written in KL1 which would not run without a PIM.

Thus, we wondered if there would be any interest from researchers outside of ICOT. Fortunately, the IFS server was frequently accessed from many countries from the beginning.

In the Follow-on Project, we made great efforts to complete the sequential version of KLIC quickly so that IFS users could have an environment to be able to run KL1 programs on their machines. We released 7 programs as IFS including the sequential version of KLIC in November 1993. As we expected, the KLIC was frequently transferred by many people.

In March 1995, as the final release of IFS, we will add 16 programs to IFS including the parallel version of KLIC and several interesting parallel application programs. They are actually operational on many workstations and parallel machines.

In October 1994, we started up ICOT's World Wide Web server [1]. We are providing general information on ICOT's activities, ICOT publications such as technical reports and ICOT journals, and outlines of major research results. ICOT Free Software is also accessible from this server.

Since its start up, the "ICOT Home Page" file has been transferred to more than 800 sites. This is very encouraging. We plan to put many more recent achievements on the server.

By using the Internet, we were able to distribute ICOT Free Software and other technical information quite effectively. Although face to face communication is indispensable for successful collaborations, we naturally rely more on computer networks for sharing advanced technical information, research tools, and products.

We plan to organize a team of researchers and engineers to maintain ICOT Free Software and the other information for a few year after the ICOT Research Center closes in March 1995.

[1] Internet addresses: anonymous FTP server for IFS is ftp.icot.or.jp, ICOT's World Wide Web server is www.icot.or.jp

4 Final remarks

The general goal of the Follow-on Project has two aspects; One is to disseminate the achievements of the FGCS Project. The other is to make a "soft-landing" for the FGCS Project which has been Japan's largest national project in the area of computer technology.

We are sure that we can attain our goal for dissemination.

First of all, the development of KLIC is proceeding smoothly. The KLIC sequential version is being smoothly ported to Unix-based workstations by many KLIC users outside of ICOT. The number of KLIC users is increasing rapidly both in Japan and other countries. They use KLIC mainly for research and educational purposes.

The KLIC parallel version is now being used at ICOT and is being debugged. The KLIC parallel version is also being ported to several parallel machines at several different sites.

Porting is rather easy. However, to get the best performance, optimization using dedicated hardware functions for parallel machines is necessary.

We expect that this problem will be solved because middle-level software which supports parallel processing such as PVM is now rapidly being optimized on many parallel machines.

Improvement and porting of the application systems are still underway. One of the problems we have in this porting is that large scale MIMD parallel machines have not yet become popular so that many researchers do not have easy access. Particularly the number of element procesasors and memory capacity is not large enough for us to port our application systems.

The expectations of users of the parallel machines seems to be very large. However, parallel machine vendors feel that the market is still small. We understand that this gap is caused by the lack of an efficient parallel programming language and environment. We are confident that KLIC is a promising solution to fill this gap.

We will also accomplish the goal of a soft-landing for the FGCS Project. We are glad that ICOT researchers are getting better opportunities and better positions in academia and industray.

Many of the ICOT graduates who returned to their companies are engaged in interesting jobs such as development of parallel machines and distributed software. ICOT researchers who were at ICOT for several years and led the FGCS Project are welcomed by many universities.

We are happy that ICOT's PIMs have been welcomed by several universities. They will be moved to these universities and used to foster researchers for the next decades.

To maintain and further disseminate ICOT Free Software, we plan to organize a virtual research laboratory on the Internet comprising ICOT graduates and collaborative research partners. This will be a loosely-coupled research group who share a common interest in ICOT Free Software. We hope that this group will act as a think-tank to foster future research projects.

References

[Chikayama 1994] T. Chikayama, "Parallel Basic Software", Proceedings of the International Symposium on Fifth Generation Computer Systems 1994, Dec. 1994.

[Nitta et al. 1994] K. Nitta, K. Yokota, A. Aiba and M. Ishikawa, "Knowledge Processing Software", Proceedings of the International Symposium on Fifth Generation Computer Systems 1994, Dec. 1994.

[DPBCT 1994] Committee for Development and Promotion of Basic Computer Technology, "Fifth Generation Computer Systems Project Final Evaluation Report", ICOT Journal, No. 40, pp-2-24, 1994.

[Uchida et al. 1993] S. Uchida, T. Chikayama, and K. Nitta, "Knowledge Information Processing by Highly Parallel Proccessing", ICOT Technical Report TR-0854, Sep. 1993.

[Uchida et al. 1993] S. Uchida, R. Hasegawa, K, Yokota, T. Chikayama, K. Nitta and A. Aiba, "Outline of the FGCS Follow-on Project", New Generation Computing, Vol. 11, No. 2, OHMSHA and Springer-Verlag, 1993.

[Uchida (Ed.) 1992] S. Uchida (Ed.), "Proceedings of the FGCS PROJECT EVALUATION WORKSHOP", ICOT Technical Memo TM-1216, Sep. 1992.