Keynote Speech

Launching the New Era

Kazuhiro Fuchi
Director
ICOT Research Center

Thank you for coming to FGCS'92. As you know, we have been conducting a ten-year research project on fifth generation computer systems. Today is the tenth anniversary of the founding of our research center, making it exactly ten years since our project actually started.

The first objective of this international conference is to show what we have accomplished in our research during these ten years.

Another objective of this conference is to offer an opportunity for researchers to present the results of advanced research related to Fifth Generation Computer Systems and to exchange ideas. A variety of innovative studies, in addition to our own, are in progress in many parts of the world, addressing the future of computers and information processing technologies.

I constantly use the phrase "Parallel Inference" as the keywords to simply and precisely describe the technological goal of this project. Our hypothesis is that parallel inference technology will provide the core for those new technologies in the future—technologies that will be able to go beyond the framework of conventional computer technologies.

During these ten years I have tried to explain this idea whenever I have had the chance. One obvious reason why I have repeated the same thing so many times is that I wish its importance to be recognized by the public. However, I have another, less obvious, reason.

When this project started, an exaggerated image of the project was engendered, which seems to persist even now. For example, some people believed that we were trying, in this project, to solve in a mere ten years some of the most difficult problems in the field of artificial intelligence (AI), or to create a machine translation system equipped with the same capabilities as humans.

In those days, we had to face criticism, based upon that false image, that it was a reckless project trying to tackle impossible goals. Now we see criticism, from inside and outside the country, that the project has failed because it has been unable to realize those grand goals.

The reason why such an image was born appears to have something to do with FGCS'81—a conference we held one year before the project began. At that conference we discussed many different dreams and concepts. The substance of those discussions was reported as sensational news all over the world.

A vision with such ambitious goals, however, can never be materialized as a real project in its original form. Even if a project is started in accordance with the original form, it cannot be managed and operated within the framework of an effective research scheme. Actually, our plans had become much more modest by the time the project was launched.

For example, the development of application systems, such as a machine translation system, was removed from the list of goals. It is impossible to complete a highly intelligent system in ten years. A preliminary stage is required to enhance basic studies and to reform computer technology itself. We decided that we should focus our efforts on these foundational tasks. Another reason is that, at that time in Japan, some private companies had already begun to develop pragmatic, low-level machine-translation systems independently and in competition with each other.

Most of the research topics related to pattern recognition were also eliminated, because a national project called "Pattern Information Processing" had already been conducted by the Ministry of International Trade and Industry for ten years. We also found that the stage of the research did not match our own.

We thus deliberately eliminated most research topics covered by Pattern Information Processing from the scope of our FGCS project. However, those topics themselves are very important and thus remain major topics for research. They may become a main theme of another national project of Japan in the future.

Does all this mean that FGCS'81 was deceptive? I do not think so. First, in those days, a pessimistic outlook predominated concerning the future development of technological research. For example, there was a general trend that reserch into artificial intelligence would be of no practical use. In that sort of situation, there was considerable value in maintaining a positive attitude toward the future of technological research—whether this meant ten years or fifty. I believe that this was the very reason why we received remarkable reactions, both positive and negative, from the public.

The second reason is that the key concept of Parallel Inference was presented in a clear-cut form at FGCS'81. Let me show you a diagram (Fig. 1). This diagram is the one I used for my speech at FGCS'81, and is now a sort of "ancient document". Its draft was completed in 1980, but I had come up with the basic idea four years earlier. After discussing the concept with my colleagues for four years, I finally completed this diagram.

Here, you can clearly see our concept that our goal should be a "Parallel Inference Machine". We wanted to create an inference machine, starting with study on a variety of parallel architectures. For this purpose, research into a new language was necessary. We wanted to develop a 5G-kernel language—what we now call KL1. The diagram includes these hopes of ours.

The upper part of the diagram shows the research infrastructure. A personal inference machine or workstation for research purposes should be created, as well as a chip for the machine. We expected that the chip would be useful for our goal. The computer network should be consolidated to support the infrastructure. The software aspects are shown in the bottom part of the diagram. Starting with the study on software engineering and AI, we wanted to build a framework for high-level symbol processing, which should be used to achieve our goal. This is the concept I presented at the FGCS'81 conference.

I would appreciate it if you would compare this diagram with our plan and the results of the final stage of this project, when Deputy Director Kurozumi shows you them later. I would like you to compare the original structure conceived 12 years ago and the present results of the project so that you can appreciate what has been accomplished and criticize what is lacking or what was immature in the original idea.

Some people tend to make more of the conclusions drawn by a committee than the concepts and beliefs of an individual. It may sound a little bit beside point, but I have heard that there is a proverb in the West that goes, "The horse designed by a committee will turn out to be a camel".

The preparatory committee for this project had a series of enthusiastic discussions for three years before the project's launching. I thought that they were doing an exceptional job as a committee. Although the committee's work was great, however, I must say that the plan became a camel. It seems that their enthusiasm created some extra humps as well. Let me say in passing that some people seem to adhere to those humps. I am surprised that there is still such a so-called bureaucratic view even among academic people and journalists.

This is not the first time I have expressed this opinion of mine about the goal of the project. I have, at least in Japanese, been declaring it in public for the past ten years. I think I could have been discharged at any time had my opinion been inappropriate.

As the person in charge of this project, I have pushed forward with the lines of Parallel Inference based upon my own beliefs. Although I have been criticized as still being too ambitious, I have always been prepared to take responsibility for that.

Since the project is a national project, it goes without saying that it should not be controlled by one person. I have had many discussions with a variety of people for more than ten years. Fortunately, the idea of the project has not remained just a personal belief but has become a common belief shared by the many researchers and research leaders involved in the project.

Assuming that this project has proved to be successful, as I believe it has, this fact is probably the biggest reason for its success. For a research project to be successful, it needs to be favored by good external conditions. But the most important thing is that the research group involved has a common belief and a common will to reach its goals. I have been very fortunate to be able to realize and experience this over the past ten years.

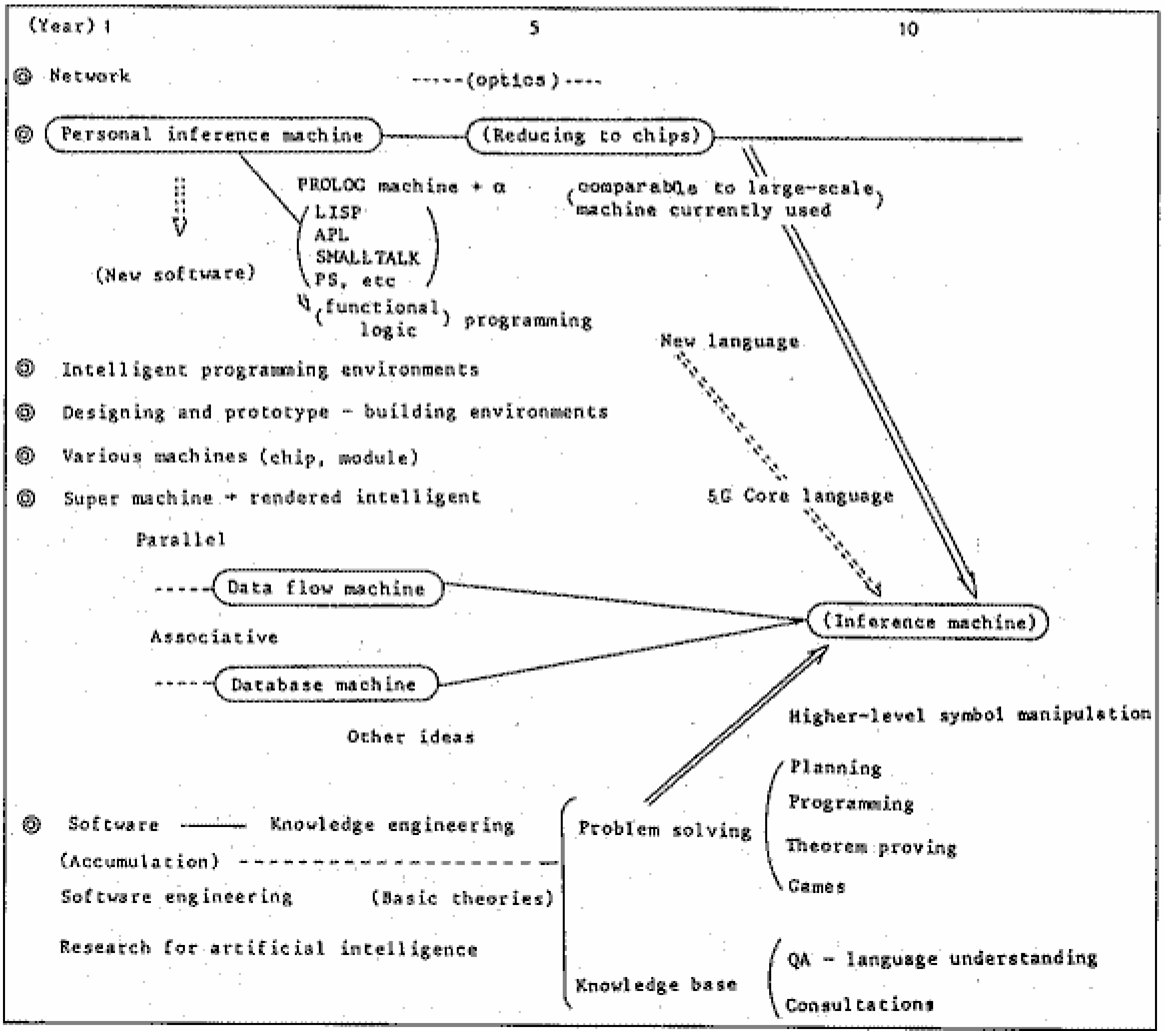

Fig. 1 Conceptional development diagram

T. Moto-oka (ed.): Fifth Generation Computer Systems (Proc. FGCS'81), JIPDEC: North-Holland, 1982, p. 113

So much for introductory remarks. I wish to outline, in terms of Parallel Inference, the results of our work conducted over these ten years. I believe that the remarkable feature of this project is that it focused upon one language and, based upon that language, experimented with the development of hardware and software on a large scale.

From the beginning, we envisaged that we would take logic programming and give it a role as a link that connects highly parallel machine architecture and the problems concerning applications and software. Our mission was to find a programming language for Parallel Inference.

A research group led by Deputy Director Furukawa was responsible for this work. As a result of their efforts, Ueda came up with a language model, GHC, at the beginning of the intermediate stage of the project. The two main precursors of it were Parlog and Concurrent Prolog. He enhanced and simplified them to make this model. Based upon GHC, Chikayama designed a programming language called KL1.

KL1, a language derived from the logic programming concept, provided a basis for the latter half of our project. Thus, all of our research plans in the final stage were integrated under a single language, KL1.

For example, we developed a hardware system, the Multi-PSI, at the end of the intermediate stage, and demonstrated it at FGCS'88. After the conference we made copies and have used them as the infrastructure for software research.

In the final stage, we made a few PIM prototypes, a Parallel Inference Machine that has been one of our final research goals on the hardware side. These prototypes are being demonstrated at this conference.

Each prototype has a different architecture in its interconnection network and so forth, and the architecture itself is a subject of research. Viewed from the outside, however, all of them are KL1 machines.

Manager of Research Department, Uchida and Laboratory Chief, Taki will show you details on PIM later. What I want to emphasize here is that all of these prototypes are designed, down to the level of internal chip, with the assumption that KL1, a language that could be categorized as a very high-level language, is a "machine language".

On the software side as well, our research topics were integrated under the KL1 language. All the application software, as well as the basic software such as operating systems, were to be written in KL1.

We demonstrated an operating system called PIMOS at FGCS'88, which was the first operating system software written in KL1. It was immature at that time, but has been improved since then. The full-fledged version of PIMOS now securely backs the demonstrations being shown at this conference.

Details will later be given by Laboratory Chief Chikayama, but I wish to emphasize that not only have we succeeded in writing software as complicated and huge as an operating system entirely in KL1, but we have also proved through our own experience that KL1 is much more appropriate than conventional languages for writing system software such as operating systems.

One of the major challenges in the final stage was to demonstrate that KL1 is effective not only for basic software, such as operating systems and language implementations, but also for a variety of applications. As Laboratory Chief Nitta will report later, we have been able to demonstrate the effectiveness of KL1 for various applications including LSI-CAD, genetic analysis, and legal reasoning. These application systems address issues in the real world and have a virtually practical scale. But, again, what I wish to emphasize here is that the objective of those developments has been to demonstrate the effectiveness of Parallel Inference.

In fact, it was in the initial stage of our project that we first tried the approach of developing a project around one particular language. The technology was at the level of sequential processing, and we adopted ESP, an expanded version of Prolog, as a basis.

Assuming that ESP could play a role of KL0, our kernel language for sequential processing, a Personal Sequential Inference machine, called PSI, was designed as hardware. We decided to use the PSI machine as a workstation for our research. Some 500 PSIs, including modified versions, have so far been produced and used in the project.

SIMPOS, the operating system designed for PSI, is written solely in ESP. In those days, this was one of the largest programs written in a logic programming language.

Up to the intermediate stage of the project, we used PSI and SIMPOS as the infrastructure to conduct research on expert systems and natural language processing.

This kind of approach is indeed the dream of researchers, but some of you may be skeptical about our approach. Our project, though conducted on a large scale, is still considered basic research. Accordingly, it is supposed to be conducted in a free, unrestrained atmosphere so as to bring about innovative results. Some of you may wonder whether the policy of centering around one particular language restrains the freedom and diversity of research.

But this policy is also based upon my, or our, philosophy. I believe that research is a process of "assuming and verifying hypotheses". If this is true, the hypotheses must be as pure and clear as possible. If not, you cannot be sure of what you are trying to verify.

A practical system itself could include compromise or, to put it differently, frexibility to accommodate various needs. However, in a research project, the hypotheses must be clear and verifiable. Compromises and the like could be considered after basic research results have been obtained. This has been my policy from the very beginning, and that is the reason why I took a rather controversial or provocative approach.

We had a strong belief that our hypothesis of focusing on Parallel Inference and KL1 had sufficient scope for a world of rich and free research. Even if the hypothesis acted as a constraint, we believed that it would act as a creative constraint.

I would be a liar if I was to say that there was no resistance among our researchers when we decided upon the above policy. KL1 and parallel processing were a completely new world to everyone. It required a lot of courage to plunge headlong into this new world. But once the psychological barrier was overcome, the researchers set out to create new parallel programming techniques one after another.

People may not feel like using new programming languages such as KL1. Using established languages and systems only, or a kind of conservatism, seems to be the major trend today. In order to make a breakthrough into the future, however, we need a challenging and adventuring spirit. I think we have carried out our experiment with such a spirit throughout the ten-year project.

Among the many other results we obtained in the final stage was a fast theorem-proving system, or a prover. Details will be given in Laboratory Chief Hasegawa's report, but I think that this research will lead to the resurrection of theorem-proving research.

Conventionally, research into theorem proving by computers has been criticized by many mathematicians who insisted that only toy examples could be dealt with. However,very recently, we were able to solve a problem labelled by mathematicians as an 'open problem' using our prover, as a result of collaborative research with the Australian National University.

The applications of our prover is not limited to mathematical theorem proving: it is also being used as the inference engine of our legal reasoning system. Thus, our prover is being used in the mathematics world on one hand, and the legal world on the other.

The research on programming languages has not ended with KL1. For example, a constraint logic programming language called GDCC has been developed as a higher-level language than KL1. We also have a language called Quixote.

From the beginning of this projet, I have advocated the idea of integrating three types of languages—logic, functional, and object-oriented—and of integrating the worlds of programming and of databases. This idea has been materialized in the Quixote language; it can be called a deductive object-oriented database language.

Another language, CIL, was developed by Mukai in the study of natural language processing. CIL is a semantics representation language designed to be able to deal with situation theory. Quixote incorporates CIL in a natural form and therefore has the characteristics of a semantics representation language. As a whole, it shows one possible future form of knowledge representation languages.

More details on Quixote, along with the development of a distributed parallel database management system, Kappa-P, will be given by Laboratory Chief Yokota.

Thus far I have outlined, albeit briefly, the final results of our ten-year project. Recalling what I envisaged ten years ago and what I have dreamed and hoped would materialize for 15 years, I believe that we have achieved as much as or more than what I expected, and I am quite satisfied.

Naturally, a national project is not performed for mere self-satisfaction. The original goal of this project was to create the core of next-generation computer technologies. Various elemental technologies are needed for future computers and information processing. Although it is impossible for this project alone to provide all of those technologies, we are proud to be able to say that we have created the core part, or at least provided an instance of it.

The results of this project, however, cannot be commercialized as soon as the project is finished, which is exactly why it was conducted as a national project. I estimate that it takes us another five years, which could be called a period for the "maturation of the technologies", for our results to actully take root in society. I had this prospect in mind when this project started ten years ago, and have kept declaring it in public right up until today. Now the project is nearing its end, but my idea is still the same.

There is often a gap of ten or twenty years between the basic research stage of a technology and the day it appears in the business world. Good examples are UNIX, C, and RISC, which has become popular in the current trend toward downsizing. They appear to be up-to-date in the business world, but research on them has been conducted for many years. The frank opinions of the researchers involved will be that industry has finally caught up with their research.

There is thus a substantial time lag between basic research and commercialization. Our project, from its very outset, set an eye on technologies for the far distant future. Today, the movement toward parallel computers is gaining momentum worldwide as a technology leading into the future. However, skepticism was dominant ten years ago. The situation was not very different even five years ago. When we tried to shift our focus on parallel processing after the initial stage of the project, there was a strong opinion that a parallel computer was not possible and that we should give it up and be happy with the successful results obtained in the initial stage.

In spite of the skepticism about parallel computers that still remains, the trend seems to be changing drastically. Thanks to constant progress in semiconductor technology, it is now becoming easier to connect five hundred, a thousand, or even more processor chips, as far as hardware technology is concerned.

Currently, the parallel computers that most people are interested in are supercomputers for scientific computation. The ideas there tend to still be vague regarding the software aspects. Nevertheless, a new age is dawning.

The software problem might not be too serious as long as scientific computation deals only with simple, scaled-up matrix calculations, but it will certainly become serious in the future. Now suppose this problem has been solved and we can nicely deal with all the aspects of large-scale problems with complicated overall structures. Then, we would have something like a general-purpose capability that is not limited to scientific computation. We might then be able to replace the mainframe computers we are using now.

The scenario mentioned above is one possibility leading to a new type of mainframe computer in the future. One could start by connecting a number of processor chips and face enormous difficulties with parallel software.

However, he or she could alternatively start by considering what technologies will be required in the future, and I suspect that the answer should be the Parallel Inference technology which we have been pursuing.

I am not going to press the above view upon you. However, I anticipate that if anybody starts research without knowing our ideas, or under a philosophy that he or she believes is quite different from ours, after many twists and turns that person will reach more or less the same concept as ours—

possibly with small differences such as different terminology. In other words, my opinion is that there are not so many different essential technologies.

It may be valuable for researchers to struggle through a process of research independently from what has already been done, finally to find that they have followed the same course as somebody else. But a more efficient approach would be to build upon what has been done in this FGCS project and devote energy to moving forward from that point. I believe the results of this project will provide important insights for researchers who want to pursue general-purpose parallel computers.

This project will be finished at the end of this year. As for "maturation of the Parallel Inference technology", I think we will need a new form research activities. There is a concept called "distributed cooperative computing" in the field of cmputation models. I expect that, in a similar spirit, the seeds generated in this project will spread both inside and outside the country and sprout in many different parts of the world.

For this to be realized, the results of this project must be freely accessible and available worldwide. In the software area, for example, this means that it is essential to disclose all our accomplishments including the source codes and to make them "international common public assets".

MITI Minister Watanabe and the Director General of the Bureau announced the policy that the results of our project could be utilized throughout the world. Enormous effort must have been made to formulate such a policy. I find it very impressive.

We have tried to encourage international collaboration for ten years in this project. As a result, we have enjoyed opportunities to exchange ideas with many researchers involved in advanced studies in various parts of the world. They have given us much support and cooperation, without which this project could not have been completed.

In that regard, and also considering that this is a Japanese national project that aims at making a contribution, though it may only be small, toward the future of mankind, we believe that we are responsible for leaving our research accomplishments as a legacy to future generations and to the international community in a most suitable form. This is now realized, and I believe it is an important springboard for the future.

Although this project is about to end, the end is just another starting point. The advancement of computers and information processing technologies is closely related to the future of human society. Social thought, ideologies, and social systems that fail to recognize its significance will perish as we have seen in recent world history. We must advance into a new age now. To launch a new age, I fervently hope that the circle of those who share our passion for a bright future will continue to expand.

Thank you.