Evaluation Report

on the FGCS Project

Wolfgang Bibel
Technical University Darmstadt
Germany

4 June 1992

Summary

In this report I try to briefly evaluate the results of the Fifth Generation Computer Systems (FGCS) project. First, I describe my interactions with the researchers involved in the project in order to make the background of my judgment visible to the reader. Second, I clarify the criteria under which I undertake the evaluation. Within the evaluation itself given in section 3 my main points of appraisal of the technical achievements are

- the demonstration of logic as a uniform and efficient framework for designing machines and software at the same time;
- the enormous gains demonstrated through using parallelism; and
- the demonstration of the gains in efficiency by producing software in logic.

Considering these and many other achievements of the project I judge it as an outstanding success. In section 4 I then test some of the major hypotheses underlying the project and come to the conclusion that all of them were solid and proved successful. In the final section I have a few remarks for the future of ICOT and of research in the spirit of FGCS which include the suggestion to continue ICOT for a limited period of time and set up a Japanese Research Institute for basic research in Information technology.

1. My interactions with ICOT research

I am proud to be able to say that I have been in contact with the key persons of the FGCS project since its preparatory phase or, more precisely, since IJCAI-79 held in Tokyo in August 1979. In discussions with Dr. Fuchi during that conference it became clear to me that he and I shared the same vision of logic offering the potential for a uniform and comprehensive approach to building and using computers for programming, problem solving, and knowledge engineering.

In 1981, I was given the privilege to present one of the six invited lectures at the first FGCS conference. In this lecture I outlined my view of software development from a logical point of view, a view that is now beginning to emerge also within the FGCS project among the applications of tools such as MGTP to program synthesis (viz. the MENDELS ZONE system).

I visited ICOT early in 1990 for about two weeks and became even more intimately familiar with many facets of the projects carried out within the framework of the FGCS project. At that time the plans for the work in the final stage of the project were just finalized which gave me a unique chance to encourage those involved in the planning task to emphasize deduction and automated theorem proving as one of the promising application areas for the basic software planned to be operative on PIMs at the end of the project.

Not only did I visit Japan several times in these past thirteen years but also had I a chance to host a number of visitors from ICOT during that time period. Dr. Fuchi and the late Prof. Moto-Oka visited the Technical University in Munich in September 1981 in order to brief me for the conference. Afterwards several other researchers from ICOT (at least ten) visited my laboratory in Munich and later the one in Darmstadt, both in Germany, for an intensive exchange of results, experiences, and opinions. Among them are Dr. Furukawa, Dr. Hasegawa, Dr. Fujita, and others.

A particular extensive exchange of our respective work became possible by a German-Japanese workshop on deduction held in 1991 at the GMD in Birlinghofen, Germany, for which I was in charge of the coordination. Eight researchers from Japan (with a large proportion from ICOT) and about twenty researchers from Germany participated.

In addition there were many occasions for encounters with researchers from ICOT at conferences such as IJCAI, AAAI, Logic Programming, Automated Deduction, etc.

All this is meant to show that my evaluation of the results finally achieved in this project is based on a rather intimate knowledge of what was going on in the project over its entire life-span. It should also inform the reader that I followed this project with a great interest and sympathy from the very beginning. In this respect my judgment may be regarded as a biased one. On the other hand, what is wrong with sharing similar visions in science?

2. Evaluation criteria

There are various possible ways of judging a project the size of the FGCS project. Depending on which of these ways one applies one would get different evaluation results. In order to avoid any possible misunderstandings, I first want to clarify which among the following ways I prefer in the present context.

a. The *accountant's* way of judgment would be to go back to the original report of the project published in 1981 and, taking it as a checklist, find out the percentage of the targets that are now actually achieved. I strongly believe that in a basic research project the size of FGCS this kind of evaluation would be rather meaningless. Therefore I will not follow this way in this report. I might nevertheless mention that according to my feeling the project has indeed achieved all its core milestone targets (while some less central topics for good reasons have been dropped along the way).

b. The *journalist's* way would judge the success against the expectations which were generated in the public through whatever process or events. The FGCS project has indeed generated various different expectations depending which public we are looking at. For instance, the Japanese press understood it very differently in comparison with the press in the US. While the public opinion is important for a project this size (since public money is involved in substantial amounts), I do not think I should bother with this issue here as a technical reviewer. Just as an aside I mention that the FGCS project has now a less favorable press in the US and in other parts of the world due to exaggerated expectations associated with the project and caused by complex reasons of a political nature in a broad sense. This includes the fact that the initial FGCS report for political reasons contained a vision of knowledge engineering which by the public was misunderstood as the final targets of the project.

c. The *economist's* way would judge the success by the amount of economic impact caused by the results of the project. Again this way does not make sense for a basic research project for which the impact is to be expected not before many more years. In this respect the project would actually have to be judged a total failure since at present the economic impact is probably close to zero. In the long range, however, it might (and probably will) turn out to be enormous.

d. The viewpoint I take is that of a *scientist*.. It consists in estimating what the net-effect of the enterprise might be, i.e. the effect of the enterprise in comparison with the situation, were it not been undertaken. As effect I understand all the changes caused by the project including the scientific results, the technology evolved, the systems and machines built, but also the changes caused in the Japanese and the international research community, or in the entire world for that matter. In addition to such an estimation I speculate about whether the net effect could have been improved by changing some of the project's conditions.

3. The project's net effect

The project has produced results and effects of very different kinds. A predominent effect is political and social in nature; others are of an infrastructural kind; and of course there are the scientific results in form of publications, systems and machines. I will discuss them all in this order.

As far as I know, the FGCS conference in 1981 was the first conference held in Japan which attracted the worldwide interest at such a high level of international visibility. For the first time the world got the feeling that Japan is about to take the lead in one of the key technologies of the future. Clearly, these feelings were mixed with serious concerns. Some people overreacted and spoke even of a technological war. Today some people again overreact. As they see that their fears have not materialized, they regard the project as a failure.

On balance, I judge the political net effect as a success. Japan has indeed proved that it has the vision to take a lead for the rest of the world. On the other hand, it acted wisely and offered the results to the international public for free use, thus acting as a leader to the benefit of mankind and not only for its own self-interest. One must, however, be aware that politically the views have not settled down to a stable state of opinions. False political steps in the future could easily destroy the current positive state of affairs.

Socially the effect of Japan's initiative is that the rest of the world has recognized the importance of information technology for the well-being of mankind. The existence of major institutions and projects in the US (MCC and others), in Europe (ECRC, SICS, ESPRIT, Alvey, and others), and in other parts of the world is to be seen as a direct consequence of the FGCS project. They all have contributed to the advancement of information technology.

One of the major results and successes of the FGCS project is its effect on the infrastructure of Japanese research and development in information technology. By an extremely clever arrangement hundreds of young Japanese researchers in research institutions of industry or universities have actively learned about the latest state of the art in information technology. This is because of the many links of ICOT with companies and universities and its policy to exchange researchers in its laboratory after their temporary stay at ICOT. Not only have these researchers learned more than would have been possible by mere education, but they also were exposed to international cooperation and now enjoy the possibility to continue these contacts at their respective institutions. Since before this project Japan had some problems with opening up to the international research community, I regard this effect as one of extreme importance for the future prospects of Japanese ability to remain a leader in information technology scientifically as well as economically. As a German I wished my country would have taken similarly wise moves in this respect, especially in the area of machine design and architecture.

Not only has the project changed the infrastructure in Japan, but also the one of the international research community. While previously western scientists rarely took their Japanese colleagues into serious consideration, now Japanese scientists in information technology are considered as equal partners a par with any others. Japanese researchers present their results more than ever before in international journals and conferences. Vice versa, Japanese journals (like the Future Generation Computer Systems Journal) and Japanese conferences (like the FGCS) are regarded as esteemed stages for the presentation of scientific results for scientists from all-over the world. The fact that Japan will host again in 1997 one of the most influential and largest conferences in information technology, namely IJCAI, underlines the respect with which our Japanese colleagues. are regarded by the rest of the world.

Finally, and most importantly, I am genuinely impressed by the scientific achievements of this remarkable project. For the first time in our field, there is a uniform approach to both hardware and software design through a single language, viz. KL1.

On the one hand, the machines built under the framework named PIM all are designed for the special purpose of executing KL1 programs which makes this execution remarkably efficient. On the other hand, all software is built on top of KL1. This is an exciting achievement for a number of reasons, some of which I will mention in the sequel.

Remember that KL1 is (sort of) a logical language. The rest of the computing world ignored logic as useful vehicle for computation mainly for two reasons, namely for its alleged inappropriateness for state-dependent software (such as an operating system) and for its inefficiency. The FGCS project has given proof that both concerns are actually wrong. Firstly, the kernel of the operating system for the PIM machines is part of KL1's realization, while the rest of the operating system is built as a large software system, called PIMOS, which is all written in KL1 using the kernel operating system functions contained in it (with about 133K lines of code). Logic can well be used as a formalism to cope with systems which are state sensitive as PIMOS proves. Secondly, the realization of KL1 is extremely efficient as the application software systems (like MGTP and many others) demonstrate in a remarkable way.

The other part of the basic software built on top of KL1 is a knowledge base management system, Kappa-P, on top of which Quixote, a knowledge representation language is built. It is less surprising that a logical language like KL1 is suitable for knowledge representation. The remarkable feature, however, is that the basis is exactly the same as the one for the operating system. The optimization efforts could therefore be concentrated on the realization of KL1 on the machines with the benefits for PIMOS and Kappa-P falling out for free.

Logic as a uniform and efficient framework is thus one of the outstanding results of the project. Aspects of this are
- the view of hardware and software design as an integral part of the problem of

information processing as a whole and
- the equal importance of inference and knowledge for knowledge processing.

The other major result is the *importance of parallelism*.. Since a lot of software produced during the project was first coded in a sequential way, the speed-up by parallelization could be experienced in an explicit and dramatic way, i.e. near linear speed-up could be experienced in a number of cases. This was by far not happen to happen, so that the international research community is grateful to the Japanese researchers to carrying out this important experiment and achieve this encouraging result. It is parallelism which eventually enabled the project to meet the performance target of 100MLIPS (logical inferences per second) for execution of KL1.

A third major point in my judgment is the ease of logic as a formalism for *efficient production of reliable software*. It is nearly unbelievable how much software was produced in about two and a half years written directly or indirectly in KL1. As one could see in the demonstrations no problems arose running these large systems. In order to appreciate this achievement in a fair way, one has to keep in mind that all this software is written for parallel execution. We all know how hard it is to code parallel programs, and in fact I know of no project anywhere in the world which has produced parallel software at such a large scale. Given the experience with conventional software production (even sequential, let alone parallel one) which obviously requires much more time for producing software with the same functionality, it is obvious at least to me that one of the results of the project is a proof for the claim that software production is enhanced by logic by orders of magnitude.

In addition to these and many other important main results, there are obviously the many results of detail, available in many hundreds of published papers and operative systems. Whatever the exact number is, we all know from our daily scientic work how many of the results of the Japanese colleagues play an important role in our own research which would not be the case without the FGCS project.

4. Evaluation of the projects hypotheses

One might speculate whether the net results of the project could have been even better, would different routes be followed, a topic which I discuss briefly in the present section.

First of all, betting exclusively on *logic* has been a real bargain in all respects as the discussed results demonstrate. The same is true for dealing with the problem in a *vertically integral way*, from hardware all along through to intelligent functions and programs.

Some people argue that it has been a mistake to test the approach based on parallel logic only at such a late stage in the project. On the one hand, there is a point to this argument because so far the computing community became hardly interested in the details of the approach simply because they could be impressed only by attractive applications. On the other hand, how could one manage to demonstrate the taken approach without having completed the machines and the basic software? I think this is a shortsighted argument. It is one of the major virtues of the Japanese way of carrying out this project that such long-range goals were undertaken and kept unchanged for such a relatively long period.

Another issue of possible concern is the specialized nature of the PIM machines, built especially to run KL1 efficiently. Would not general purpose parallel machines (like the J-machine presented in an invited lecture at the conference) serve the same, if not a better purpose? I think this is a good question which cannot be answered at present in a fully satisfactory way. I only can speculate about the outcome of future experiments focussing on exactly this issue. It would be extremely surprising to me, if specialization would not make a difference, rather I believe it does make a difference. Especially for the implementation of a logical language (rather than an imperative or functional one) it may be crucial to use hardware specialized to carry out the required logical operations. At present the progress in machine development might still outway this advantage: until a specialized machine is ready for operation, the general purpose machines have become so advanced in performance that they easily compensate the disadvantage of being unspecialized. I am convinced, however, that ultimately machines dedicated to the style of programming (logical, functional, imperative) will become crucial, especially for logic programming. So the experiment carried out within the FGCS project will turn out to be extremely valuable for future logic-based machine design. In appreciating this judgment one should keep in mind what I said in the previous section concerning the value of logic in producing reliable software fast, since people tend to forget the investment in terms of people's time spent for software production in machine comparisons.

As to the choice of KL1 there might be the concern that it is not really a logical language inpure form. Although this is in fact true, I still believe that KL1 is the best compromise which could be achieved at the time of the design of this particular language. I do expect a new and improved approach in this respect in some future project; but it is the fate of any project that at some point in time its results will be improved by further scientific progress.

5. Perspectives for the future

Given the outstanding success of the FGCS project I think it would be a waste of efforts, would the valuable infrastructure of ICOT and the basis in terms of machinery and software achieved now be abandoned. In other words, I strongly recommend that ICOT shall continue to exist for alimited period of time (e.g. five years) in some form or another. Its tasks might be to

- evaluate the machines and software systems w.r.t. their crucial features;
- exploit the results in various applications;
- maintain the systems; and
- pursue appropriate new research goals.

Maintenance of the systems is especially important in view of the fact that MITI has adopted the policy to make all software available as free software which is a remarkable step which will have its effect in terms of international cooperation. Of course, this policy will bear its fruit only when the software will become available on standard machines other than PIMs; but as I understand there is anyway the plan to port it to a UNIX environment.

There will be changes in the personell currently leading ICOT. I want to emphasize the enormous influence on the success of the project which is due to its scientific director Dr. Fuchi. Although I fully understand his desire to take a rest for his own personal research goals, it would be a real pity if one would not use his talents for some other, perhaps even grander enterprise.

With having achieved such a high reputation for carrying out basic research in a successful way, the idea of a scientific institution for basic research based in Japan occurred to me. It could continue to play part of the role currently played by ICOT to maintain close links with researchers from all over the world. It could be a meeting place for first-rate researchers from all over the world. As an aside I mention that institutions of this nature would best be placed in an environmentally attractive area.

As I indicated in the previous section, I strongly believe that with KL1 an important, but not final step was taken towards a purely logical machine and software. New directions in logic (such as linear logic) will have to be taken into account for doing a next step in the same direction. In any case, I am convinced that the problems with conventional software production will bring the rest of computer scientists eventually towards the same line of a logically oriented computing and knowledge engineering discipline.

Résumé

Wolfgang Bibel is Professor for Intellectics at the Department of Computer Science of the Technical University Darmstadt in Germany where he heads a group of more than ten researchers. He also maintains affiliations with the University of British Columbia as Adjunct Professor and with the Canadian Institute for Advanced Research as Associate.

In 1968 he received his Ph.D. degree from the Ludwig-Maximilian University of Munich, Germany. For many years he worked at the Technical University of Munich as a Senior Researcher, building up the AI group there. In 1987 he became a Professor in Computer Science at the University of British Columbia, Vancouver, and a Fellow of the Canadian Institute for Advanced Research.

His more than one hundred publications range over various areas in artificial intelligence such as automated deduction, machine architecture for deductive systems, program synthesis, knowledge representation, but include also topics concerning the implications of AI technology for society.

Dr. Bibel is Section Editor of the Artificial Intelligence Journal, Associate Editor of the Journal for Symbolic Computation, Co-Editor of the AI book series of Vieweg Verlag, and on the board of more than ten further journals and series. From 1982 through 1986 he served as the first Chairman of ECCAI, the European AI organization. Since 1987 he is a Trustee of the International Joint Conferences for Artificial Intelligence, Inc., and held the Conference Chair of IJCAI'89. In 1990 he has been awarded the title of an *AAAI Fellow* by the American Association for Artificial Intelligence (AAAI) "in Recognition of Significant Contributions to the Field of Artificial Intelligence".