Report on FGCS Project

H. Gallaire

Surely, it is difficult for me to report on a project with which I have had many points in common. ECRC, the laboratory of Bull, ICL, and Siemens which I had the role to direct in its formation years had a lot to share with ICOT, in spirit if not in practice. I hope that my comments will not be distorted by what I would like others to think about ECRC's own role, achievements and problems.

First I want to command the ICOT management for getting external viewpoints, for opening wide their doors, and for letting so much of their work be discussed.

1 SUMMARY

The scientific success of the FGCS project is not to be disputed. There has been innovative work, there has been deep and thorough work yielding a better understanding of issues which would otherwise not have progressed as much.

The technical achievements of ICOT are impressive. Given the novelty of the approaches, the lack of background, the difficulties to be solved, the amount of work done which has delivered something of interest is purely amazing; this is true in hardware as well as in software. The number of alternatives explored must have generated a large body of knowledge which, if it can be shared will prove to be a major return for all of us. The steps taken to provide free software for all is definitely a positive one, even if it is currently hampered by the absence of hardware available to run it or by lack of specifications available to all, as appears to be the case for e.g. KL1. Whether many people will take advantage of this offer remains to be seen because it appears to be difficult to take only bits and pieces of it.

The scientific vision of the ICOT management and of their sponsors has been maintained all over these years, and it must have been a very challenging task; they must be complimented for holding firm to their technical vision. The enthusiasm behind this vision has had positive effect on the whole community across the world, leading others to take actions, if not to follow the approach that ICOT was pursuing.

The fulfillment of the vision, should I say working on the 'grand plan' and bringing benefits to the Society, is definitely not at the level that some people anticipated when the project was launched. This is not, to me, a surprise at all, i.e. I have never believed that very significant parts of this grand plan could be successfully tackled.

I was expecting however to see "actual use" of some of the technology at the end of the project. There are three ways in which this could have happened. The first way would have been to have real world applications, in user terms: only little of that can be seen at this stage, even though the efforts to develop demonstrators are not be be underestimated. The second would have been to the benefit of computer systems themselves (eg impacting the computer manufacturers); this does not appear to be directly happening, at least not now and this is disappointing if only because the Japanese manufacturers have been involved in the FGCS project, at least as providers of human resources and as subcontractors; whether this lack is due to the fact that not enough effort went into getting their true support (which may have been a tough issue after TRON and SIGMA) or not, is difficult to assess from a distance, and the responsibility for that state of affairs certainly lie in several hands. However, I firmly believe that when Japanese industry starts looking for engineers and designers for parallel systems (which may already be the case), they will draw heavily on the skills developed at ICOT in the FGCS project; indeed what has been learned through the development of many versions of machines, of parallel OS, of cache management and of load distribution algorithms, of distribution networks, etc will undoubtedly be useful to them. The third way would have been to impact computer science outside of the direct field in which this research takes place: for example to impact AI, to impact software engineering, etc; not a lot can yet be seen, but there are promising signs, eg the results on AI in legal reasoning or theorem proving; by the way there are again direct ways and indirect ways through which the project impacts these fields: by making sheer use of the powerful hardware technology and making practical what was known but was impractical on conventional hardware (the parallel theorem prover is one such example), or by true innovation using the new tools of the project ; there are certainly several examples of the latter (eg CAL, QUIXOTE, ...); it seems to me however that there has been more reliance on the use of the power of parallelism; this is probably natural since developing parallel systems was and still is the major technical thrust of the project; one can only wonder what are the limits of this position, as we were reminded by A. Robinson quoting M. Minsky during his invited talk at FGCS'92. More application work would have been needed to feel fully optimistic about the impact on the environment of this work.

The project made a choice of one approach of symbolic computing, namely logic programming (LP); it pursued it very consistently; this is a very wise behavior, and I did the same at ECRC. ICOT went very far, building many different pieces of hardware (convincing us if needed of the exceptional manner in which technology is mastered in Japan), building full operating systems with great success, investigating many solutions in parallel. If one wants to establish a new center, I would recommend to follow the same pattern, namely stick to one type of technology, especially when it is a new technology and when so little is known about it. This allows one to see a problem under different yet related views and helps progressing to solve it. For example the work on parallel implementation of KL1 is useful in its own right but it also provides insight for other problems as well, e.g. for KBMS implementation.

I believe logic is the right choice that had to be made to investigate knowledge based systems; I have argued elsewhere that it does not mean that there is only logic in the practical world but there is no contradiction here. Perhaps this perspective should have guided ICOT more towards integration of logic to other environments than it did.

I will not criticise the choice of KL1 against the choice of logic languages a la Prolog; it is important to understand the limits of each approach; if the parallel implementations of Prolog a la Andorra work well (on truly large scale problems), fine; otherwise we know we have solutions a la KL1; it is early in the game to know for sure.

In general I am surprised positively by the speed at which the researchers have picked up the background that they encountered elsewhere during their research to make novel proposals; constraint logic programming is one such example where the progress made is significant, even though they were not the first players.

When it comes to discuss specific results, it is difficult to single out one of them, because the areas covered are so different. I only would like to mention again the fact that all the work done on parallel systems implementations will definitely bear fruit in a non distant future, directly or indirectly. I feel that the work on knowledge bases is not as foolproof as some of the other work done in the project and that the QUIXOTE environment, although its is appealing when one considers all the features it integrates, would need more testbench work before it can be adopted because I find it complex and lacking some of the features that such complex semantic representation systems need (see below). The work on parallel theorem provers embodies some nice results and has shown that it goes beyond state of the art; however I have some reservation due to the fact that parallelism cannot be the answer to all difficult problems. The work on constraint languages is very interesting and one of the very few to allow to use non linear constraints; this work shows the high level of skill with which the developments have taken place; there is also room for improvement here because again speed is probably not the only answer. I will not comment much further individual results except to say that case based reasoning may appear to be easier to do now than before (until we run into other speed bareers ...)

2 FUTURE WORK

There is a list of actions which could be mentioned here; I will only stress some of my main points.

There is a need to evaluate how standard technology will support the FGCS results, knowing that the standard technology itself is not stable, in terms of performance, in terms of features (distributed OS may appear soon for example). This is in my mind crucial to the future of the results of this project.

I believe that there is a need for ICOT to show the impact of the technology on classical problems. Just to give an example, why not try to develop a payroll package, where knowledge bases and rules could play an essential role for building easy to customise packages, where parallelism is clearly possible and interesting at the processing and database level, where constraints can be useful for a human resource package (eg for allocating people to tasks, ...) and where it is possible to enhance basic packages by AI: for example finding the best person qualifying for a task, based on a description of skills etc; this can be as sophisticated as one wants, but only as the icing on the cake. I am sure there other potential applications where such combinations are possible. What would be important would be to be able to compare development time, maintenance time, adaptation time, performance and costs of running systems, etc. In general I feel a need for more comparative work, taking into account costs which I admit may be difficult to do since the new hardware and software cost can hardly be compared to commercial one.

The results on parallel implemetations are impressive; however there is the need to work on automating the mapping between programs (in KL1) and processors; if this is not done, it may jeopardise building the higher level languages and applications which need to run efficiently in most cases and to exploit well the architectures.

There is a need to work on important issues related to the knowledge bases work, at least to simplify it and to address for example the problem of integrity constraints (which is not, as I understand it, what has been done under the constraints heading in QUIXOTE). I also believe that the work on constraints need more research and incorporation of more propagation-like techniques. I have no feeling about what's needed for the work on the Genome project.

There are definitely enough results obtained and enough good and important problems waiting for an answer, that there is no doubt that a follow up of the project is needed.

Hervé Gallaire - 2 June 92

Hervé Gallaire
GSI
25 Boulevard de l'Amiral Bruix
75782 Paris Cedex 16