Energy, Mines and Resources Canada
Énergie, Mines et Ressources Canada

CANMET
Canada Centre for Mineral and Energy Technology
Centre canadien de la technologie des minéraux et de l'énergie

PROCEEDINGS OF THE CANMET SYMPOSIUM ON EXPERT SYSTEMS IN THE MINERAL INDUSTRY

OTTAWA, ONTARIO

CAMSELL HALL,
588 BOOTH STREET,
OTTAWA

MARCH 31, 1987

Edited by D. Laguitton

SP 87-4E

Available in Canada through

Associated Bookstores
and other booksellers

or by mail from

Canadian Government Publishing Centre
Supply and Services Canada
Ottawa, Canada K1A 0S9

Catalogue No. M38-15/87-4E — Canada: $25.00
ISBN 0-660-12651-6 — Other Countries: $30.00

Price subject to change without notice.

FOREWORD

This manual contains the transcripts of the presentations made at the CANMET Symposium on Expert Systems, on March 31, 1987. The transcripts were produced from the sound recording of the sessions except for the last presentation for which only the visual aids used by the author are reproduced. Whenever available, the key visual aids of other papers have also been included in their original form. The reader must therefore bear with colloquial sentences and expressions. Distributed as a kit with the videotape entitled "The Computer Age in Mineral Processing", these proceedings will reinforce the underlying message that in our industry as in others, obsolescence and bankruptcy can be fought today by an early integration of the technology that will prevail tomorrow.

D. Laguitton

TABLE OF CONTENTS

SESSION 2

Video Tape Display of the Texas Instruments Presentation Kit on Expert Systems: how to select a project; how to start. First and Second Artificial Intelligence Satellite Symposium recordings. Texas Instruments, P.O. Box 181153 Austin, Texas 78718

OPENING ADDRESS

DR. K. WHITHAM, ADM,
RESEARCH AND TECHNOLOGY SECTOR,
ENERGY, MINES AND RESOURCES CANADA

It is with great pleasure that I accepted the invitation to open this CANMET symposium on Expert Systems, and to welcome you to EMR from all parts of our country. We are very pleased that we have representatives here from industry, university and government, and that industry is represented both by users of technology and by the suppliers of technology and related services.

Rarely have we seen a more typical example of the rapidly evolving technology that is characterizing the late twentieth century, than with that of Artificial Intelligence (AI) and more specifically with Expert Systems which attempt to imitate the way humans reason and solve problems.

As the technology suppliers and related services are striving for a share of a contemplated enormous market, the potential users have often only a weak notion of the potential of such advanced techniques, and a vague comprehension of the terminology that describes the proposed new tools. Such terms as fuzzy logic, knowledge based systems, inference engines, frames and rules, object programming, forward and backward chaining, heuristics and shells need to be explained to the user to allow informed choices between the many options offered in software solutions to his or her process engineering challenges. Demonstrating tools and prototype applications are very powerful means of upgrading technological awareness in a target industry. The department of Energy, Mines and Resources is proud to be able to lead such a technology transfer exercise today, in the form of a full day symposium on expert systems in the mineral processing industry.

CANMET has a long tradition of technological innovation and technology transfer. Scientific contributions to the field of mineral process simulation can be traced back to papers and reports in the early seventies when computer penetration was still minimal in the country's concentrators and even in mineral processing corporations at large. The theme became more visible in the late seventies and early eighties as computer power became distributed to remote terminals through data networks. At this time CANMET chose to launch its SPOC project on Simulated Processing of Ore or Coal. The objective was to inject a first generation of simulation software into the newly available hardware stations deployed in Canadian mineral processing industries. Contributions were sought from the many centres of excellence in the field from universities, government and private sector research centres, to consolidate and exhaustive base of methods and software for process evaluation and optimization by computers. Computers programs and associated documentation were produced over a period of 6 years to assist process engineers in their daily process monitoring, to increase the accuracy of plant audits, design and retrofitting, to facilitate equipment selection and design or to expand teaching and training of mineral processing techniques. Technology transfer was activated from the very early days of the SPOC project and, by 1985, some 70 organizations had already elected to acquire mainframe software and documentation from CANMET. This number is now over 130. The impact of such a transfer became visible in published engineering studies acknowledging the use of SPOC software, and in numerous industrial implementations of selected software acquired from CANMET. International recognition came along, as a number of request and positive feedback came from every major mineral producing country.

Computer technology, meanwhile, has continued its fast innovative pace, and by the mid 80's, the microcomputer market had pretty much settled around PC compatible standards. CANMET was prompt to seize this opportunity for easier software transfer, and the entire collection of SPOC Fortran programs was converted to be delivered on PC diskettes. Besides the distribution of software and manuals, technology transfer was implemented by several applied workshops covering specific domains of computer applications in mineral processing. About one hundred key individuals in Canadian companies have availed themselves of these training opportunities.

Every engineer knows that the notion of transfer implies an emitter, a relay and a receiver. CANMET has fulfilled its mission as emitter of new methodology or as relay of methodology emitted in the country's numerous mineral engineering research centres. Workshops and seminars have also been used to promote the required adjustments in the "receiving" industry. A steering committee composed of industrial members had been established to ensure the industrial awareness and monitoring of the SPOC project. But the necessary adjustments in the industrial workforce have been slower than the technological innovation. Penetration of the computer technology is still largely controlled by the natural renewal of the working force, as new, computer literate staff enters the field.

It is very clear that Technology Transfer in the field of process simulation is more than a question of exchanging diskettes or instruction manuals. However the situation may change more rapidly in the future as we see the arrival of expert systems that will play an important role in the transfer of expertise, including computer expertise. As a logical extension of the SPOC project, the Mineral Sciences Laboratories of CANMET have therefore undertaken to integrate expert systems as R&D tools for Computer Aided Mineral Processing, under the acronym SPEX, The Simulation of Process EXpertise. The power of the expert system would be increased over that of numerical methods by the capture of the numerous rules of thumb and other heuristic knowledge developed by experienced operators and engineers throughout the years. Using pseudo natural languages, the expert system promises to be a very powerful training and diagnosis tool, and will soon be integrated into on-line process control software.

As a preamble to the 1987-1988 activities in this area, MSL has sponsored two pilot studies of feasibility that aim at assessing development tools for typical expert systems, one based on conventional architecture microcomputers, the second based on a state of the art LISP machine and a powerful expert system shell. These two prototypes will be discussed in more detail during the symposium. One of these applications is done in the leaching plant of a zinc refinery and the other applies to a clinker grinding circuit in a cement operation.

CANMET is therefore already able to provide guidelines on how to access this new technology and this symposium is the first real opportunity for potential users in the mineral industry to get acquainted with expert systems methodology as well as to establish contacts with the suppliers of AI tools.

With this technological bridge in mind, I am therefore pleased to wish you a most successful symposium. I invite you to take some time to discuss your different engineering problems and seek innovative solutions that will capture old, perishable but valuable knowledge into permanent and upgradable expert systems, to establish contact with staff in my department that may lead to successful technology transfer, and most of all, establish communication channels by which CANMET can keep aware of your particular problems and challenges, successes and failures. In this way we can continue to meet tomorrow's challenges in the Canadian mineral processing industry.

AI AND EXPERT SYSTEMS: CONCEPTS AND DEFINITIONS

D. LAGUITTON, RESEARCH SCIENTIST, CANMET

My objective in this presentation on expert systems is to define the terms and to avoid that every following presentation consist in an introduction of expert systems. You know this is the most introduced subject of the year 1986 and will also most likely be the most introduced subject in 1987 but we are in front of an audience who is more interested in demonstrations than in introductions, therefore, I will deal with the introduction and I hope that other speakers will restrict themselves to demonstrations.

We have in this audience two broad families of people representing two lines of activities, the first group is represented by people who are active in artificial intelligence. They are not all salesmen but as we know, the most visible ones are salesmen. Besides the salesmen we have also people who have been active during the last 30 years in developing the tools and the theory of artificial intelligence as a science. Now we are almost run over by the AI salesman who comes with deliverable systems and abbreviations such as AI, KBS, LISP, shells, these are the terms I would like to define for those of you who may not be familiar with them. For those of you who may know them better than I do, I will just say, come and see me later on and tell me where I was wrong. The other group of people we have in this audience is represented by the mineral processing engineer who thinks in very square terms and who wants solutions. His job is to produce commodities and not to decorate his process with the latest christmas tree of computer technology, hardware or software. If that equipment can help him do his job better he will use it but he is pragmatic and what he wants is solutions. Between these two different worlds, these two different specialties, there is a wide gap that can be bridged by technology transfer. Technology transfer is as you know or as you may have learned during the introductory presentation of Dr. Whitham, and the following video projection, one of the missions of CANMET. Technology transfer consists in making people of both worlds meet, like today and we hope you will seize the opportunity to meet between representatives of providers of AI technology and representatives of users of AI technology. We are also trying to support your joint undertakings in our projects at CANMET. One of the characteristics of an expert in any domain of specialization is very often to consider that anyone who does not belong to his type of activity knows nothing about his field. I have a word of wisdom for this type of individual and this word is that "there has been an alarming increase in the number of things I know nothing about". I say "I" because this applies to me but I suspect it also applies to many people in this audience. Therefore, we have to be ready more and more to make use of technology transfer to acquire the expertise that is not in our traditional domain of training. One of the advantages of having gathered representatives of those two worlds is that providers of AI technology have been relying a lot on the concept of prototyping and this concept fits particularly well the frame of mind of people in the mineral processing industry. These people are accustomed to the concept of pilot plant to try things on a small scale to avoid that a catastrophy occurs on a larger scale, therefore, tool assessment by prototyping seems to be a very promising way for both worlds to exchange problems and solutions.

At CANMET we started being involved on a major scale with simulation in 1980 and in 1986 we have added to our activities the development of expert systems (Fig. 1). All along we keep our minds on the ultimate objective of process control. Our efforts in expert systems development have been carried along

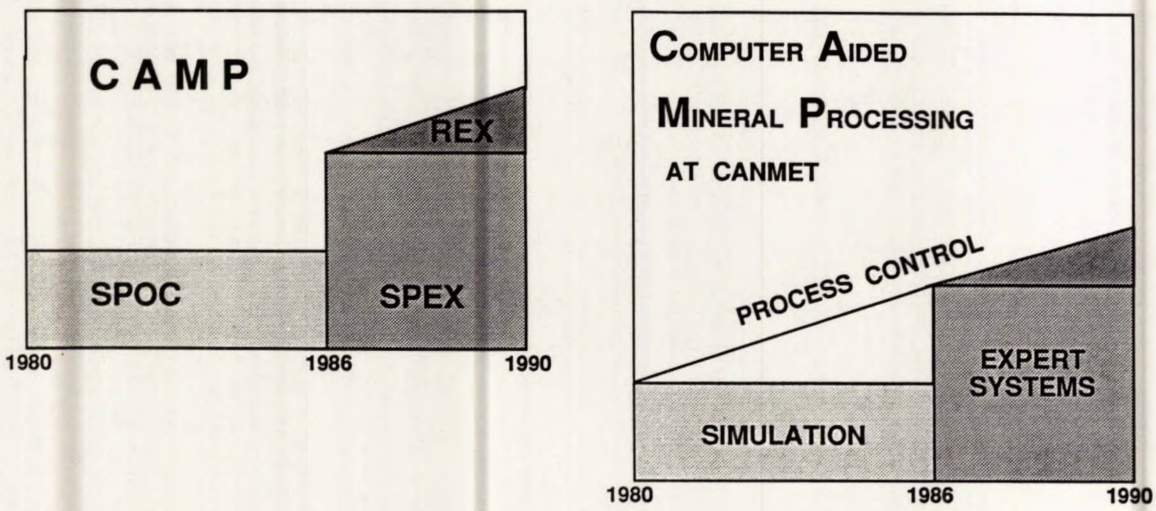

Fig. 1 Computer aided mineral processing at CANMET since 1980

AI TOOLS ASSESSMENT BY PROTOTYPING		
TOOLS:	PC	KEE / EXPLORER
DOMAIN:	ZINC	CEMENT
USER:	CANMET LAVALIN CEZ	CANMET LAFARGE UNIVERSITE LAVAL UNISYS

Fig. 2 The two approaches to expert systems prototyping sponsored by CANMET in 1986

two lines in 1986. We have chosen to assess AI tools by prototyping using two families of tools: one is based on PC computers, personal computers, and the other one is based on LISP machines, which are dedicated AI computers (Fig. 2). In the first family of PC based tools, we are evaluating an expert system shell called Personal Consultant from Texas Instruments. The domain in which we apply this expert system shell is in a zinc refinery and the team involved in assessing this tool is composed of CANMET, Lavalin, and CEZ (Canadian Electrolytic Zinc). The second family of tools under assessment at CANMET is represented by the TI Explorer LISP machine, and the KEE shell, (Knowledge Engineering Environment). The KEE software is developed by Intellicorp and this application is carried out in a cement factory. The team is composed of CANMET, Canada Lafarge Cement, Laval University, and Unisys, the company that resulted from the merging of Sperry and Burroughs. These two projects will be described in more detail in following presentations. We will also hear a presentation by Dr. Swinkels on the future of expert systems in extractive metallurgy as he perceives it. Now lets go to some definitions.

What is artificial intelligence. Well, when you tell me what intelligence is, I will tell you what artificial intelligence is. It is basically the same thing, done with a machine, it is the simulation by an artificial device, a machine, a computer, of the behaviour that is usually considered as being intelligent in a human. A definition proposed by Texas Instruments is that AI is the study of how to make computers solve problems traditionally thought to require human intelligence. The illustrations I am using for this introduction are those distributed by Texas Instruments in their excellent presentation kit that includes slides, manuals and videotapes. In the last twenty

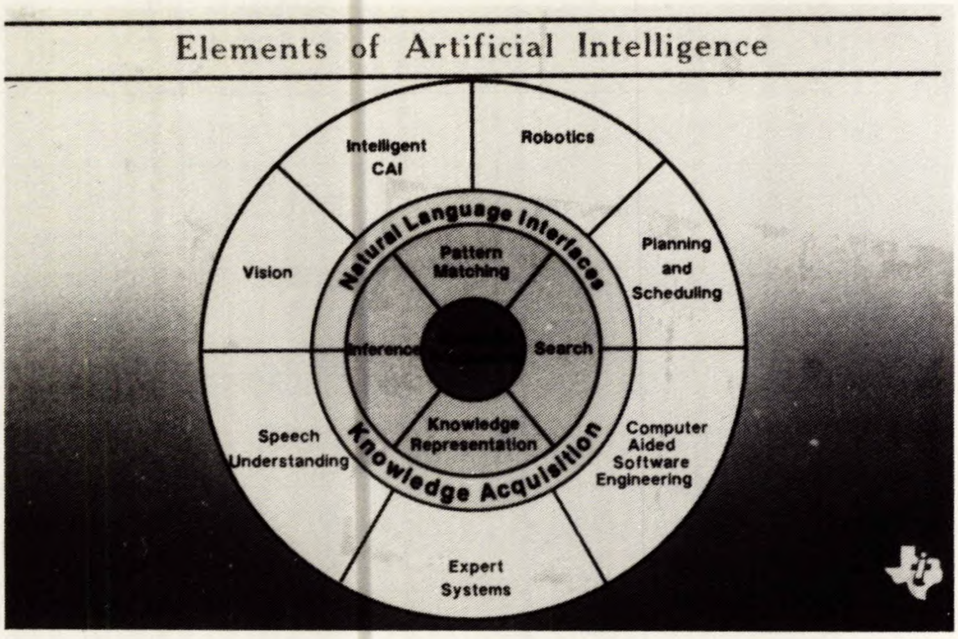

Fig. 3 The elements of Artificial Intelligence

or thirty years the world of computers has evolved from being used as big number crunchers to being used as database managers, followed by use as information processing devices for word processing or communication of data from one network to another, and now we see computers having entered the age of knowledge processing. This area of knowledge processing is precisely the field of application of expert systems. The different elements of artificial intelligence are presented on Fig. 3, at the center of AI we have symbolic processing, as opposed to numeric processing, which was a traditional field of computer applications. Symbols can be words, drawings, or any other symbolic way that humans use to represent knowledge. Symbolic processing is used to perform some function such as pattern recognition, searching for a particular solution within a finite of possible solutions, knowledge representation itself, or inference which is the word used to describe the actual reasoning process. To interface with these functions of AI, we can use natural languages, or pseudo-natural languages and knowledge acquisition methods. Finally, with these interfaces to AI, we can, as users, develop applications in robotics, in planning and scheduling, in software engineering, in expert systems, in speech recognition, artificial vision or computer aided instruction. Symbolic processing can therefore be defined as the representation and manipulation of knowledge and information encoded as symbols, in a way that simulates human reasoning. These symbols may represent objects, concepts, properties, or relationships. For example, a grinding mill is an object, a

classifier is an object, grinding or size reduction are concepts, viscosity, colour, temperature are qualities that can be represented by symbols.
Figure 4 is a reflection on the differences between conventional software development and symbolic software development. In conventional software development the user needs first to specify his requirements as well as the functionality of his system, to meet with a programmer/analyst who is going to design and draw the specifications of the particular software that is being sought and these steps are going to be followed by implementation of the software. At this stage the user may wish to bring some corrections and the whole process has to be repeated in an iterative fashion. The idea to remember

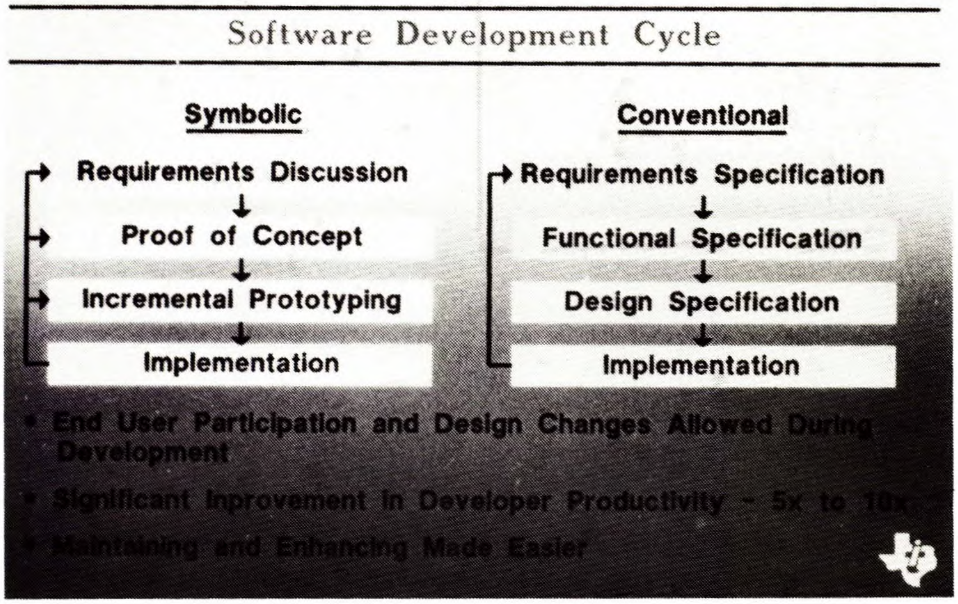

Fig. 4 Differences between conventional and symbolic software development

here is that modifying software written in conventional languages such as Fortran or others, is not a trivial task and requires familiarity not only with the language but with the overall design specifications of the system. By comparison symbolic processing can be used to develop software in a much more flexible environment, the discussion of the requirements can be followed by the proof of the concept through very rapid prototyping using only a few of the ideas or rules that are to be implemented in the final system. If the concept is valid and is accepted, development of further prototypes by incremental prototyping can be undertaken until the system reaches the level of final implementation. All the way through, the process of modifying the specifications of the functionality of the system or of the knowledge on which the software is based is much more flexible and user-friendly than it is the case with conventional software development. It is usually accepted that the main quality required to develop an expert system is to be an expert in human communication rather than be an expert in data processing. Actually being an expert in data processing may end up being more of a handicap rather than an advantage in developing expert systems.

What are expert systems? Expert systems can be defined as computer programs which use an expert's knowledge and experience to solve narrowly focused complex problems. In other words, we could define expert systems as software that simulate the art part of an expertise that fits the definition of an art and a science. Expert systems are based on the results of artificial intelligence research in the areas of knowledge representation, inference methods, explanation, and natural language. Development of an expert system requires the interaction of a domain expert, who is a person specialized in the area of knowledge for which the expert system is wanted, and a knowledge engineer who is the modern day equivalent of the programmer analyst in conventional programming. The knowledge engineer will capture the domain expertise of the expert and convert it into an expert system composed of a knowledge base, rules and facts. To do so he uses the development interface of an expert system shell so that eventually, as is often said, the expert doesn't retire but is converted into a diskette. Once the expert system is developed, the user can have access to the knowledge that is embodied in the system through the user interface and the explanation facility. The expert system can be used to produce recommendations to the user when fed by system data or to act as a consultant in reply to questions asked by the user.

A question that is often asked, is why LISP, and what is LISP? Well LISP is the symbolic processing language that has been in existence for many years and has come to light with the recent progress of artificial intelligence. It is not the only language of artificial intelligence, other languages such as PROLOG, NIAL, are also used on a large scale. LISP has become quite prevalent in North America, and this prevalence is in part due to the development of the so called LISP machines which are dedicated hardware in which many of the functions of the LISP language have been imbedded and have become system functions. LISP is used for symbolic processing. It provides great flexibility, extensibility, good interaction features and an excellent development environment. A LISP machine can be described as made of the following components: a specialized hardware, a powerful user interface, LISP language functions, as I said, imbedded in the system and a very good programming environment, which greatly facilitates the development of expert systems. As we will see, microcomputers offer also a very interesting alternative and a number of expert systems shells have been commercialized that operate on microcomputers of the PC type.

A few words now on the methods of knowledge representation. There are three basic methods, knowledge can be represented by rules, by frames, or by semantic networks which I will discuss briefly. A rule is composed of two parts, a premise and a conclusion, or an action (Fig. 5). A premise such as "IF THIS" and a conclusion such as "THEN THAT" or "THEN DO THAT". If you like to drive fast then drive a sports car, if the power draw of the grinding mill is decreasing then the circulating load of the mill is probably increasing due to laminar conditions inside the mill. A frame is the representation of an object such as for instance a car, a frame can also be called a unit or a class (Fig. 6). It has properties attached to it, such as the type, the price, and other properties as may be required. Knowledge can be represented by a tree, composed of different frames, for example, a frame called car can be subdivided into two subclasses, my car and John's car. Each of the subclasses, can

---

KNOWLEDGE REPRESENTATION - RULES

A RULE CONTAINS:

- Premise
  - The "if" part of the rule

- Action
  - The "then" part of the rule

example:

"if" likes-to-drive-fast

"then" buy-sports-car

Fig. 5 Rule representation

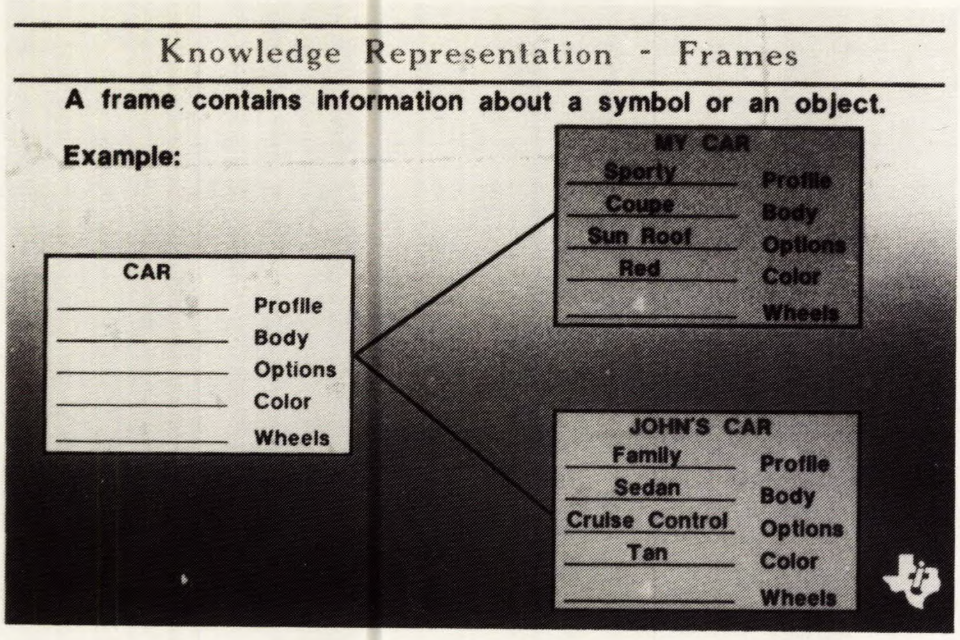

Fig. 6 Frame representation

KNOWLEDGE REPRESENTATION - SEMANTIC NETWORKS

Symbols and their relationships to each other are represented in semantic networks.

Example:

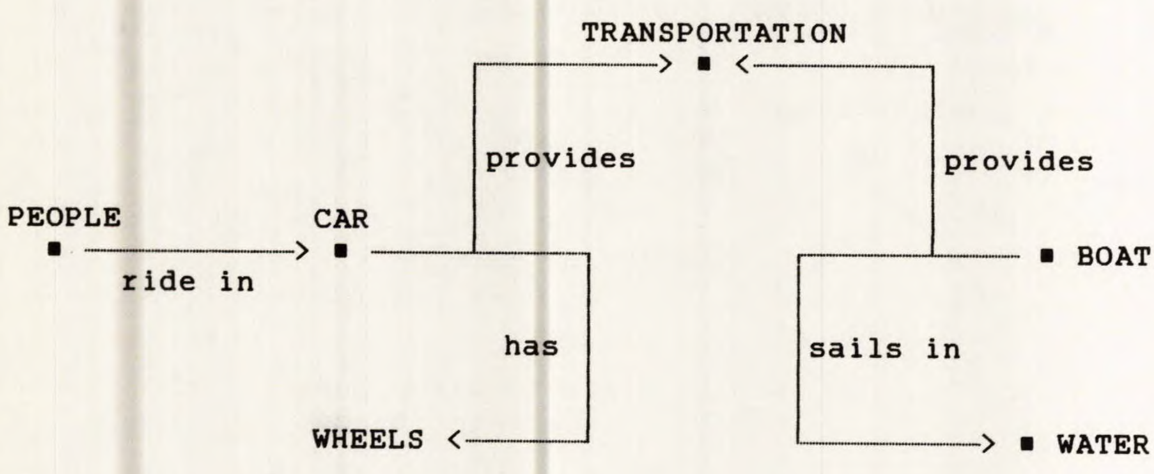

Fig. 7 Semantic network representation

have particular values for the set of properties defined in the frame. My car for instance can be a sports car and John's car can be a family car. A particular knowledge base system may control the inheritance rules of properties between the various frames represented in a tree. Some properties may be inherited by all members of the subclasses. Other properties may be locally defined for particular members. In a semantic network, the objects such as cars, boats, people, etc. are linked by semantic expressions such as "BELONGS TO", "CARRIES", "IS GREATER THAN", that define the relationships between the various objects of the tree (Fig. 7). For example, people ride in cars, a car has wheels.

Now I would like to say a word about the inference mechanisms used in expert systems. Inference is the name of the reasoning process used by the expert system. There are essentially two controlled mechanisms for inference. One is called forward chaining and the other backward chaining. Forward chaining or data driven reasoning is the type of reasoning activated when some information is provided about data contained in the premise of a rule. For instance, if one has the following rules: "IF A THEN B", and "IF B THEN C", whenever the statement "A IS TRUE" is provided to the inference engine, the activation of the rules in a forward manner results in the conclusion "C IS TRUE". This conclusion is reached by forward chaining. In some cases the number of possible outcomes or conclusions from an initial statement may be very large. This is known as combinatorial explosion. To avoid this, backward chaining or goal driven reasoning can be used. In that type of reasoning, the triggering mechanism is a statement belonging to the conclusion part of the rule. To resume the example used above, "IF A THEN B", "IF B THEN C", backward chaining will be triggered by selecting conclusion C and asking the system to verify if C IS TRUE. By backward chaining from the second rule the system will verify whether B IS TRUE, and in order to assess if B IS TRUE, the system will have to verify whether A IS TRUE. If A IS TRUE then B IS TRUE, if B IS TRUE then C IS TRUE. The system can therefore reply "yes" or "no" to the query depending on the value of A. Expert system technology provides flexibility, allows the processing and the capture of knowledge whether this knowledge is public (for instance, knowledge contained in textbooks or handbooks), or private, (for instance the knowledge possessed by a particular expert). Expert systems allows also to distribute expertise throughout an organization or to bring more uniformity to the level of expertize throughout the different shifts of the daily operation. The expertise of the senior expert can be made available to newly employed staff, not only allowing them to acquire experience at a faster pace but also allowing them to make well informed decisions earlier without the assistance of the actual human expert.

When is an expert system appropriate? The first criterion as mentioned earlier, is that the application must belong to a domain of narrow focus. The second criterion to consider is whether there are only a few individuals who have shown ability in solving problems in the domain being considered. Of course, if you have a thousand people who can do the job and you don't mind paying them, why use an expert system to perform this particular task. The third criterion is whether an expert is actually available and willing to help in developing the expert system. This is very important since an expert system is essentially a copy of the knowledge of an expert, you have to have access to the original in order to be able to make the desirable copy. Then you have to consider whether the problem to be solved is significant, is clearly definable, and is of intermediate difficulty. It is not advisable to start your first application of expert system in an area of major difficulty. Further criteria are for instance that the problem has to be appropriate for solving with symbolic processing. If your problem can be solved by straightforward mathematical modelling you do not need an expert system. Do you have commitments of management for your project? This particular point could be debated since the advantage of a successful prototype may be actually to convince management to invest further into AI applications. Current users of expert systems can be found in very diverse areas, such as, diagnostic, configuration, scheduling, monitoring, analysis and interpretation of data, consulting by user, planning, for instance planning of a project or planning of a budget, designing. Some expert systems are already linked to computer aided design software in architecture or engineering applications. Finally, expert systems work very well for providing training, instruction, and explanations.

How to get started in AI is a topic which is well covered in the video tapes of the first and second symposia on expert systems distributed by Texas Instruments. In a nutshell, you have first to identify the problem and identify the people who will be trained in order to solve this problem using AI techniques. You have to develop a first prototype of an expert system in order to prove the concept and, as the concept is proven and supported by management, you have to continue enhancing the system by integrating further knowledge and improving other features of the expert system including its functionality.

As indicated earlier, CANMET has funded in 1986/1987 two projects of tool assessment by prototyping. Without taking the fire out of the following two presentations, I just want to review briefly the conclusions that were drawn from CANMET's point of view on these two projects. In one case we had an application in a zinc refinery, the tools selected were microcomputer based, more specially an IBM PC microcomputer and the Texas Instruments Personal Consultant expert system shell. The second project involved the application of a LISP machine, the Texas Instrument Explorer, combined with the powerful KEE expert system shell distributed by Intellicorp. This second application was done in a cement grinding plant and consisted in simulating the reasoning of an expert when troubleshooting for defects in the product specifications or in the throughput of the plant. A very important conclusion that was drawn from these two applications is that you should not neglect the learning curve of the tool, a large and powerful tool such as KEE and TI Explorer require substantial basic training and several months of daily practice to be able to exploit fully the multiple capabilities of such systems. Smaller shells such as Personal Consultant or very small shells available commercially require substantially shorter training periods but of course, have much less power of inference and much fewer user facilities. To summarize, large tools benefits are that they allow for rapid prototyping, rapid modifications of the system, flexibility of inference methods, powerful user utilities. Their drawbacks are that they have a slow learning curve, they are expensive, there may be problems to deliver them on a large number of delivery systems, and there could be problems of networking as well. The small and medium size shells on the contrary, have as advantages, a fast learning curve, easy delivery, since they run on conventional hardware, easy networking for the same reason, and they are inexpensive. Their drawbacks are that they are rigid and do not provide as much flexibility in the inference method and they have limited user facilities. Detail review of these two prototype developments are going to be discussed by the following speakers and will appear in scientific literature as they reach completion.

DEVELOPMENT OF AN EXPERT SYSTEM TO ASSIST THE OPERATOR OF A MODERN ELECTROLYTIC ZINC REFINERY: PRELIMINARY RESULTS

C. GHIBU, LAVALIN MONTREAL,
AND
D. DUPUIS, CANADIAN ELECTROLYTIC ZINC,
VALLEYFIELD, QUEBEC

It is with great pleasure that I am going to communicate to you today the preliminary results of the project on pilot application of expert system technology to process control. The project has reached about 20% of its development, it started in November 1986 and is due to be completed in March 1988. The objective of this project is to develop the prototype of an expert system to assist the operator of a complex industrial process in making his decisions. We have already designed, programmed and installed several computer control systems throughout the world and we consider this project as a natural extension of our efforts in supervisory control.

BACKGROUND

Throughout the years Lavalin has been constantly at the forefront of innovations in the area of industrial computer applications. These efforts have resulted in the installation of some of the most performing control systems in the area of water treatment and cement manufacturing. We have attempted to develop and even partially achieved operator support functions that are very similar to those we are pursuing in this project and this was done in particular in a cement factory in the U.S. and in a water filtration plant in Canada. Having quickly recognized the potential of expert systems we have followed very closely their evolution. The first attempt to install an expert system has been carried out at the end of the 70's. After having designed the system we have realized that the software and the hardware that were available at the time did not have the required power. Around 1981 we have started to formalize engineering expertise to apply it to computer aided design. The main result has been the development of powerful graphic tools, of man/machine interfaces and especially the understanding of human problems associated with the implementation of a complex computer system. In 1984 with the market availability of the new expert system development tools and the considerable decrease of the price of computing power, we have reconsidered our approach. We have sent some of our experts to training workshops and we have tested in-house several expert system shells. In May 1986 we have hosted a seminar on expert systems which has been attended with considerable interest. We are presently involved in several projects using expert system technology, one of them is the subject of this presentation. Another is for the opening and closing of the mobile roof of the olympic stadium, another one for the training of operators and a fourth to estimate the cost of engineering projects.

OBJECTIVES

The main objectives of this project are,

a) evaluate the available tools including the microcomputers and the programming aids, in order to find out if they are applicable for on-line operator assistance,

b) structure a methodology for the application of expert system technology in industrial control,

c) assess the impact of this technology in the control room of a continuous complex industrial process. Please note that I am using the words technology of expert systems because we are convinced that there is no expert system as such, but only computer based systems which have access, to achieve their functionality, to expert system technology as well as to graphic animation, databanks or fuzzy logic to name just a few.

PARTICIPANTS

This project is being pursued through the efforts and expertise of several participants: CANMET, with its experience in development and transfer of mineral processing technology, CEZ: a division of Noranda, which has operating expertise as well as a strategy for assessment and implementation, the National Research Council: it participates through the ergonomy section of the systems development laboratory to help us integrate and evaluate the impact of the project on the human users. Now I would like to introduce Mr. Denis Dupuis, who is the Chief of Hydrometallurgical Operations at the plant of Canadian Electrolytic Zinc, who is going to explain the interest of his company for this project.

CEZ INTRODUCTION
BY D. DUPUIS

Canadian Electrolytic Zinc belongs to the Noranda Group. We have a capacity of producing 230,000 metric tonnes of zinc and we are the second largest producer of this commodity in the world. Our total production is 3 to 4% of the world consumption. We have a few side products, for instance we produce 430,000 tonnes of sulphuric acid and almost 500 tonnes of cadmium per year. When the plant was built more than 25 years ago, several sites have been considered and Valleyfield was selected for at least three main reasons. The first is the availability of electrical power near the canal of Beauharnois, the second is the availability of qualified manpower, we are very close to Montreal, and the proximity of a complete transportation network that includes roads, railroads, and the St. Lawrence Seaway. Now we will describe briefly our process (Fig. 1). Our raw material is a zinc concentrate, containing approximately 50% zinc, 33% sulphur and 10% iron. In a first step, the sulphides are roasted at temperatures of approximately 900° centigrade, sulphur is recovered as sulphur dioxide, converted into SO_3 and converted to sulphuric acid. Zinc is then leached into solution, the solution is chemically cleaned and then submitted to electrolysis. The metallic zinc obtained by electrolysis is then melted into ingots of different shapes and sizes. The hydrometallurgical part of our process includes lixiviation, cleaning and electrolysis. The lixiviation occurs in three successive steps, 80 to 85% of the zinc is leached during the first step, about 15% in the second step and

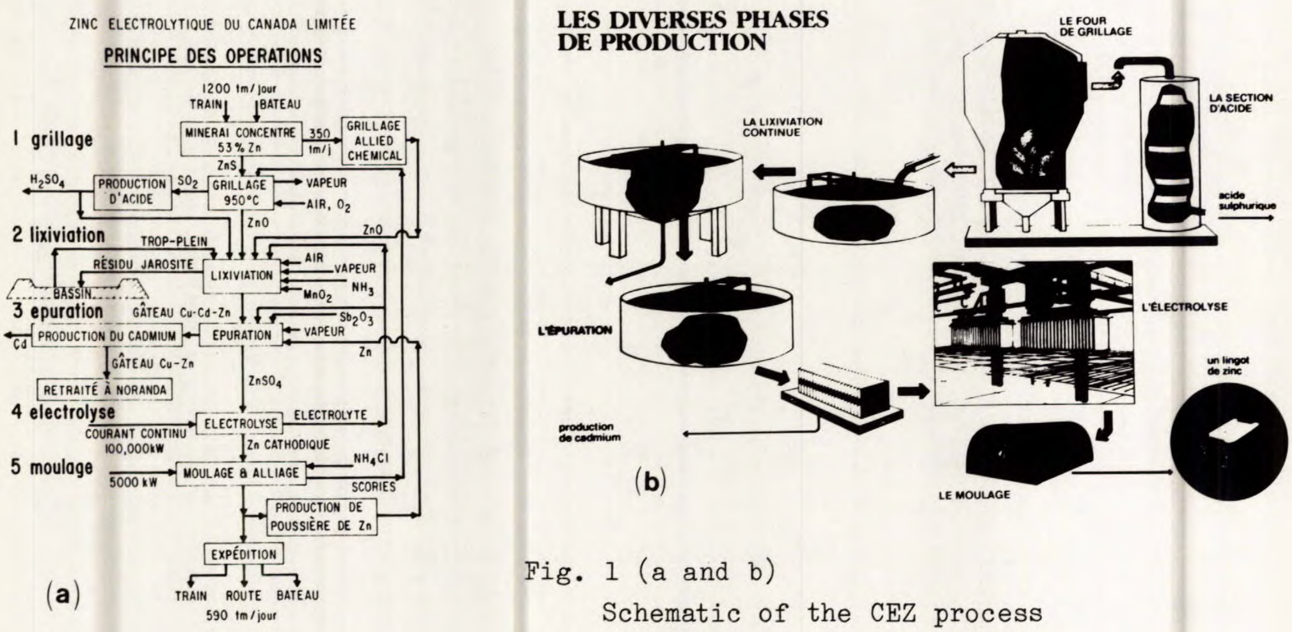

Fig. 1 (a and b)
Schematic of the CEZ process

the rest, which is about 5%, in the third step. The third step has a double role. It allows the lixiviation of zinc ferrites and to precipitate iron in the form of ammonium jarosite. The jarosite precipitate is washed in order to remove any zinc solution that may be attached to it and then it is sent to some tailing ponds. The main chemical reaction during the cleaning stage is an oxydo-reduction between metallic zinc powder and impurities that occur in ionic form in the solution. Copper, cadmium, cobalt, nickel, are cemented by the zinc powder and leave the solution. During electrolysis zinc is converted from soluble ionic form into a solid metallic form that is deposited on the aluminum cathode. During electrolysis the solution is heated as a result of the joule effect and it is necessary to cool it by vacuum or atmospheric refrigerants. I must also mention that the overall hydrometallurgical circuit can be considered as a closed circuit since the sulphuric acid generated during electrolysis is recirculated to the lixiviation stage. For zinc producers and operators of a circuit like ours, the efficiency of the hydrometallurgical process can be measured by different methods. For instance the leaching stage will be assessed by the percentage of zinc that is successfully extracted into solution. During the best periods this percentage can reach 99% while during the most inefficient periods it is down to about 97%. During the cleaning stage we want to know of course the percentage of impurities in the solution. We want this concentration to be below 1 milligram per litre for cadmium, below 0.2 milligram per litre for cobalt and 0.1 milligram per litre for copper. Electrolysis is assessed by the efficiency in terms of electrical consumption. During the best months, efficiency of electrical consumption is about 92% while it reaches down to around 87% during the worst periods. Another criterion that can be applied to the overall operation is the percentage of time at which the plant operates at full capacity. This number can be between 85 to 97 or 98%. Each of these criteria has a very high economical impact and it is necessary to maintain operation at the high levels of efficiency mentioned above on a regular basis. We have a threefold strategy in order to achieve this optimum operation. First, a year ago we started a program of operator training. They are presently attending elementary chemistry courses in order to be ready to receive specific chemical training applying to the particular section of the circuit they will be assigned to. Second, we have undertaken to introduce computers and automatic control in our process. We want as well to eliminate routine tasks such as sample collection, titration and also to speed up the transfer of process information which is currently done manually. Third, and this is the topic of the presentation today, we are proposing to install an expert system which will have the effect of improving the quality of the decisions made by the operators as well as to move to a mode of preventive maintenance rather than the curative maintenance mode in which we currently operate. Mr. C. Ghibu will now resume his presentation of the project on expert systems.

STRATEGY

Since we are dealing with an R & D project, we have selected an approach step by step in order to be able to monitor our progress and if necessary to back track without major consequences. The selected steps are,

1) conceptualization, i.e. project definition,
2) proof of concepts by a very rough prototyping in order to assess the decisions made during the conceptualization,
3) prototyping.

The conceptualization stage must answer some basic questions such as, define the functionality of the system define the user, his domaine and his level of expertise, define the user interface. Who is the expert, his domain and his level of expertise? What expertise will be extracted from other sources than from the expert? How is the system to be evaluated, what are the best tools available for developing the prototype? The main question to be answered is: what are the priorities to be considered when developing this expert system?

We have identified three main directions for the CEZ plant. The first priority is the identification of trends, that is the ability to interpret data as a function of their context. Detect the trend before this trend can lead to an upset or a production shutdown. Second priority is to diagnose problems themselves, that is to find the cause of the trend or the cause of an observed problem. The third priority is to provide the operator with a sound advice system on decisions to be made. In order to achieve these objectives we have selected an architecture as shown on Fig. 2. The first component of our system will collect the process data and produce reports that will allow the monitoring of the process as well as of the expert system. The expert system itself is composed of two independent parts, one which we call the primary diagnostic system, this part operates constantly in a closed loop and doesn't do anything else but detecting trends in the process. This is because we

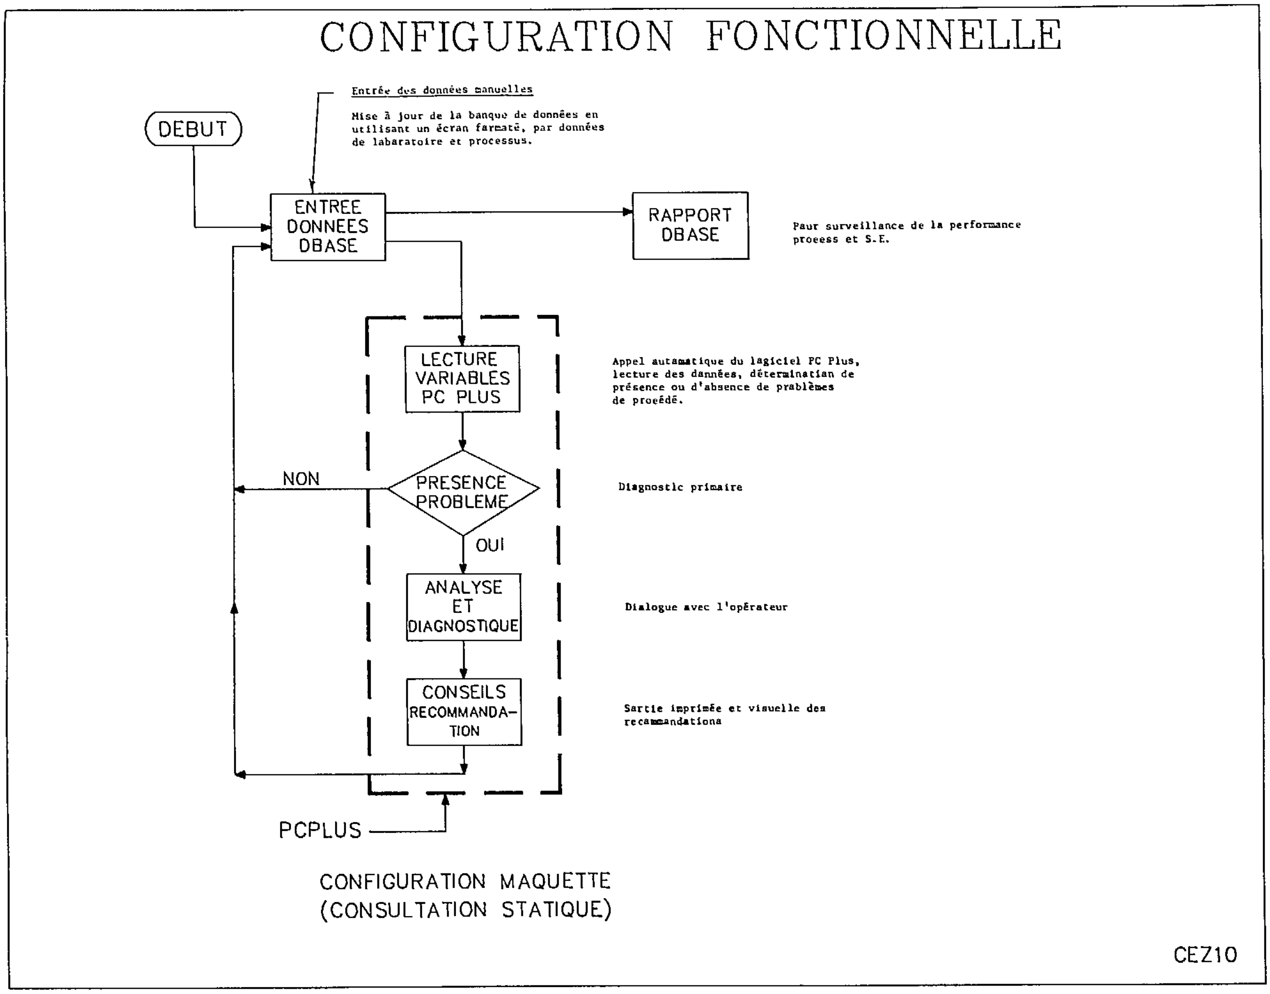

Fig. 2 Selected architecture of the CEZ prototype

have observed that the detection of trends is a key element of the operation of such a plant. If this first part of the expert system detects the trends, then it communicates with the second part, which makes a detailed diagnostic and issues advice to the operator. Following a dialogue with the operator, the expert system issues a series of advice on actions to be taken. Whether the operator follows or doesn't follow this advice is something that will be refined in advanced stages of implementation of this prototype. We strongly believe that expert systems are going to be used as part of overall control systems rather than in a stand-alone fashion, they will have to be fed with numeric or symbolic data by other components of the automatic control system and we can use values and observations which usually are not a part of conventional control methods such as the appearance of the solutions in the thickeners. This aspect can be characterized as brown or cloudy or yellow. This type of value was not traditionally used to inform the computer control system. I believe that the true originality and contribution of expert system is in making use of this information which is known by the operators to be so useful for handling the process. Another very important notion incorporated in expert systems is that any value used by the system to make a decision is attached to a degree of confidence which in turn results in a level of probability for the conclusion from the reasoning. For example, the system may warn the operator that it is likely that thickener no. 2 is under malfunction, with a probability of say 80%. Connected to this aspect we can also associate the degree of confidence in particular process parameters with the age of these parameters. In other words, the older the determination of a given variable, the less contribution it is allowed to have on a current diagnostic of the plant.

A few words now on the tools we have selected as programming aids. After evaluating a few commercially available shells we have selected Personal Consultant Plus released by Texas Instruments for the initial version of the prototype. We do not know yet whether this will also be the shell selected for the final prototype delivery. It is our opinion that the choice of the tool for the initial development is not critical since as we reach the limits of the development tools, the knowledge that has been accumulated so far can be easily transferred to another more adapted expert system shell. Therefore, contrary to what is usually claimed in commercial circles, we believe that the key element is the collection of knowledge from the expert and not the tool used to process that knowledge.

PROGRAMMING TOOLS

We have considered the following shells:

- Expert-Easy
- Insight-II
- Good for procedures and diagnostics
- We used it for the diagnostic of the programmable robots PB-400 of Merlin-Gerin

THE EXPERT

- French shell offering forward and backward chaining
- Can handle about 1000 rules on APC
- We use it to train cement plant operators

RULE-MASTER

- Tool - kit to develop expert systems specialized in diagnostic and control
- The rules are generated from examples
- Uses uncertainty factors and fuzzy logic
- Easy interface with other systems written in C

PC +

- Appropriate for the proof of concept
- Combines user-friendliness and power
- Flexible enough
- Forward and backward chaining variety of parameters
- Acceptable interface with outside world
- Some still graphics
- Vendor support and on-going upgrading
- Runs also on larger machines
- Too slow on 8087. Acceptable on 80286. Good on 80386
- Does not allow compilation
- The frames are actually nothing more then modules.

*INTELLIGENCE COMPILER

- Being evaluated. Seems promising but test is still incomplete

HARDWARE

IBM XT
ABM AT
COMPAQ 386

- ES/P - Advisor
- OPS-83
- Flops
- Rule-Master
- PC + (1)
- Intelligence-Compiler (2)

(1) Selected for the proof of concept
(2) Under investigation

DESCRIPTION OF THE SHELLS

Expert-Easy	- Easy to use. Ideal for first exposure to this technology. - Builds rules from examples - No interface with outside world - Good to structure knowledge
OPS-83	- More a language than a programming aid - Can support many different applications but requires a lot of programming
FLOPS	- Very powerful for fuzzy logic applications and handling of special confidence coefficients - can communicate with other programs for data transfer or within a procedure
M.1	- Spin-off of emycin, based on backward chaining - If ... then ... rules with uncertainty coefficients - Allows modifications of the inference method - Good user interface
ES/P ADVISOR	- Written in prolog - Dual inference (Backward chaining within paragraphs, forward between paragraphs) - Allows prolog methods - No uncertainty coefficients

HARDWARE

(XT) - 8086	Too slow for development, but can support a small application for consultation
(AT) - 80286	acceptable for development and good consultation
(COMPAQ) - 80386	Good for development. The constraint is the 640K memory addressable by DOS We are expanding our unit to 2MB in Extended Mode.

TOPICS INCLUDED IN THE EXPERTISE

- Identification of trends: interpretation of the trends as a function of context variables
- Problem diagnostic:
  - to find the cause of a trend or of an existing problem
- Advice to Operators:
  - to advise operators on recommended action

PARAMETERS

In order to activate the reasoning that leads to an identification of trends and problems, the system must use process parameters.

We have observed that these parameters are different from those used in classical control systems or in production reports.

MAIN DIFFERENCES BETWEEN CONVENTIONAL CONTROL VARIABLES

(1) Symbolic Values
    Colors
    Quality
(2) Degree of certainty
(3) Age
  - Laboratory data is used in the system
  - Operator observations play a role in reasoning
  - Rules of thumb generate state variables which are considered during the reasoning

THE PROJECT PHASES

During the proof of concept, the system is not on-line.

Process data is entered by the operator who also produces the process reports in the same operation.

During the second phase data will be directly collected by a distributed control system.

PERFORMANCE EVALUATION

During the development of the expert system it is necessary to monitor progress and to assess the quality of results. These evaluations must be done at the right time with well defined criteria and techniques. A poor evaluation method may ruin a valid development study.

A well conducted evaluation motivates the team and avoids dead-ends.

The evaluation must be explicitly planned and carried out during the entire project.

If the criteria established during the conceptualization change, everyone must be informed.

The system will be evaluated at two levels:

(1) Technical performance
(2) Acceptability in the industrial environment

Since expert systems are affecting an area which had remained exclusively under human control, some strong reactions to their implementation can be expected.

The current prototype of our expert system is on display in the commercial exhibit and we will be glad to show you a demonstration.

DEVELOPMENT OF AN EXPERT SYSTEM PROTOTYPE FOR PROCESS DIAGNOSIS IN THE CLINKER GRINDING CIRCUIT OF A CEMENT FACTORY

J. VANDERSTICHELEN, CANADA CEMENT LAFARGE, MONTREAL

Ladies and Gentlemen, it is almost a diskette that is going to speak to you since I am the expert that is being used to develop the expert system for the clinker grinding circuit. First I will ask for the projection of the short movie on the preparation of cement in order to situate the context of our project. The transcript of the sound track of this movie is inserted in these proceedings.

TRANSCRIPT OF VIDEOTAPE ON PORTLAND CEMENT

Most cement plants are located near their sources of raw materials. In this age of dwindling natural resources these fortunately are found in great abundance across the earth. Thick deposits of limestone in its many forms supply calcium the major ingredient of cement. From beds of clay, shale and slate comes the compounds of silicon, aluminum and iron vital to correct cement formulation. The materials are usually extracted from quarries or stripped from the surface and transported to the crushing plant for the first step in the process of reducing mountains to micron particles. Rock from the quarry is dumped directly into primary crushers which reduce its to sizes suitable for handling and storage. Enormous boulders are crushed into pieces the size of baseballs. To conserve energies some plants drop hot exit gases from the burning stage to the crushers and the rotary dryer when the raw materials have a high moisture content. From the primary crusher the material usually travels to a secondary crusher which further reduces it to a size generally not exceeding 2 cm. The crushed rock is then conveyed to storage areas where each materials is stored separately. In the plant laboratory the raw materials are analyzed to determine the right combination of limestone, shale, clay iron or other materials needed to produce Portland cement. They are automatically proportioned, blended and conveyed to storage for the next critical stage of manufacture which is raw grinding or preparation of the materials for burning. In the massive grinding mill the raw material is pulverized to a fine powder.

This model of a roller mill shows how heavy wheel type rollers crush the material against a rotating table. Hot gases from the kiln or clinker cooler sweep the fine particles out of a top discharge boat to a dust collector or storage silo. Raw grinding can also be done in rotating ball and tube mills. Thousands of steel balls carried up the wall of the cylindrical mill tumble back onto the material and crush it to the fineness needed for the burning process. As this model demonstrates, we have been following the materials through the dry process method of preparation. They now enter a long rotary kiln for the conversion into new mineral compounds. Cement plants use the dry process when the available raw materials have a low moisture content or other characteristics which favor dry processing. Some plants use a semi-dry process in which the raw mix if fed into a large rotating pan where water is added. The mixture forms into small dust free pellets which pass through a pre-heater before entering the kiln. Where raw materials have a high moisture content cement plants generally use the wet process. Here water is added to the materials and the grinding mill producing a creamy mixture called slurry. After grinding the slurry is sampled and pumped into blending and storage tanks. The mixture is stored in tanks under agitation until ready for the kiln. The rotary kiln is the heart of the cement making process. A long sloping steel surface lined with fire brick and with a burner at the lower end. The raw mix in the form of powder pellets or slurry is feed in at the upper end and undergoes chemical changes within the fiery chambers. It slides and tumbles its way towards the burner. First water and then carbon dioxide is driven off in the hottest zone. Calcium silicates, aluminates and other compounds are formed that give Portland cement its binding ability and strength. In the hottest stages, materials can reach a temperature of more than 2700°F or 1400°C. When the heated material emerges from the kiln in the form of red hot marble sized glass hard balls it is a new product called clinker. To save energy a number of plants are now pre-heating the raw mix before it enters the kiln. The pre-heaters are basically a series of horizontal compartments or vertical cyclone chambers through which the raw mix passes on its way to the materials through progressively hotter stages. The heat from these gases triggers the desired chemical changes so that the raw feed needs only a relatively short time in the kiln, on hour or less. Some suspension pre-heaters contain an auxiliary furnace or pre-calciner to further improve fuel efficiency. Suspended in a turbulent vortex of hot gases the raw feed releases up to 95% of its carbon dioxide. The decarbonated or calcined material is separated from the gases in stage four, then passes through a short rotary kiln where the transformation to clinker is completed in less than 1 hour. The short kiln means increased production, less radiation heat loss than long kilns and because it burns less fuel, lower fire brick replacement costs. Through the use of these highly efficient pre-heater systems, many cement plants have more than doubled their production output while dramatically reducing their fuel consumption per ton of cement. The cement industry is converting almost totally to coal as the principal kiln fuel to further conserve natural gas and oil supplies. In the central control room all phases of production from raw material crushing to kiln operation from cement storage to pollution control are monitored and automatically regulated. Computer systems keep operators instantly informed on demand, about conditions in the plant. They also provide automatic process control and operating and alarm surveillance. After the clinker leaves the kiln it may pass through a long cooler chamber where big fans force cool air through the hot clinker as it is carried along on a travelling grate. The heat removed is ducted to the kiln or pre-heater as combustion air or the cooler may be the planetary or satellite type where the red hot clinker falls into tubes mounted symmetrically on the shelves. Cooling air enters through holes in the tubes, picks up heat in the clinker and then is 100% utilized as combustion air. Pollution control device is in the form of electrostatic precipitators, batteries of fabric bag filters or gravel beds remove particulates from exit gases before they enter the atmosphere. They are found throughout the plant wherever air pollution might occur. This strict control of emissions enable the cement plant to maintain high air quality standards. After the clinker is cooled it may be removed to storage, shipped elsewhere for grinding or passed directly to the final stage of cement making. Before final grinding a small amount of gypsum is added to the clinker. The gypsum controls the setting and strength developement properties of the cement. Here in the finished grinding mill the clinker is ground to a super fine powder composed of micron sized particles as small as 1/25,000 of an inch in diameter. It can now be considered Portland cement. The cement is now so fine that it will easily pass through a sieve that will hold water. Throughout the final process the freshly ground cement is tested frequently to assure that it meets rigid quality standards. From the grinding mills the cement is conveyed or pumped to batteries of tall silos where it is stored awaiting shipment. Most bulk cement is shipped directly to customers directly by transport trucks or railroad cars. These are gravity loaded at the storage silo by overhead equipment that can fill a large tanker in only a few minutes. A small percentage of the finished product is still packaged in the familiar bags and shipped to customers who need only small amounts of cement. Semi-automated machines can fill and seal a bag in 3 seconds. Where plants are located on waterways large quantities of both bagged and bulk cement are transported by ship or barge to distribution terminals.

DESCRIPTION OF THE PROTOTYPE

Now I would like to speak about our project and the prototype we have developed. This project was started four months ago the objective was to develop a prototype of expert system for diagnostic of process defects in a clinker grinding circuit. As you have seen in the movie, clinker grinding is the last stage of cement preparation. To develop this prototype, we have used a LISP machine, called the TI Explorer, manufactured by Texas Instruments, and the expert system shell called KEE, released by Intellicorp. We have used the apprenticeship program of UNISYS to gain access to these two powerful prototyping tools. We have also relied on the contribution of the Group de Recherche sur les Applications de l'Informatique dans l'Industrie Minerale (GRAIIM) of Laval University.

The project objectives were as follows. We wanted first to assess the expert system technology for the cement industry since the cement process is rather complex, we have restricted ourselves to the rather simple area of clinker grinding and in this part of the process, we have more specifically selected the diagnosis of process defects. We wanted to assess the tool KEE/Explorer as well as the technique used for knowledge collection and representation in a specific domain. I must also mention that the Lafarge-Coppee owns over 40 cement factories and that expertise is consequently very widely distributed in this group. The centralization and the management of this expertise is therefore of the utmost importance for the company. It would be most valuable if the best expertise could be made available equally to all operations. The project, started by two weeks of formal training in the KEE and LISP languages, followed by a three month period of apprenticeship, during which we developed the prototype under supervision and advice from knowledge engineers at UNISYS.

A few words now about the KEE software which is a very powerful and very interesting piece of software. First we have access to a system of data entry that allows us to store all the static knowledge of the circuit, such as the names of all the process units with associated properties. KEE uses a frame based representation. The properties of each frame can be inherited by members of the knowledge base or conversely defined as local properties for particular objects. KEE allows also the use of knowledge represented by rules, representing the heuristics of the system under study, KEE uses an pseudo natural language for rule editing. We found this feature of KEE convenient but rather limited and in several cases, some LISP structures have to be used in order to activate particular relationships between premises and conclusions. Rules have the classic form "IF A THEN B", but we can also introduce an action as a result of the execution of a rule and this is done by entering a "DO" command. Such a rule has therefore the form, "IF A THEN B, AND DO C". One of the main drawbacks we found in KEE is the lack of an "ELSE" feature in the writing of the rules, which has forced us to write a number of

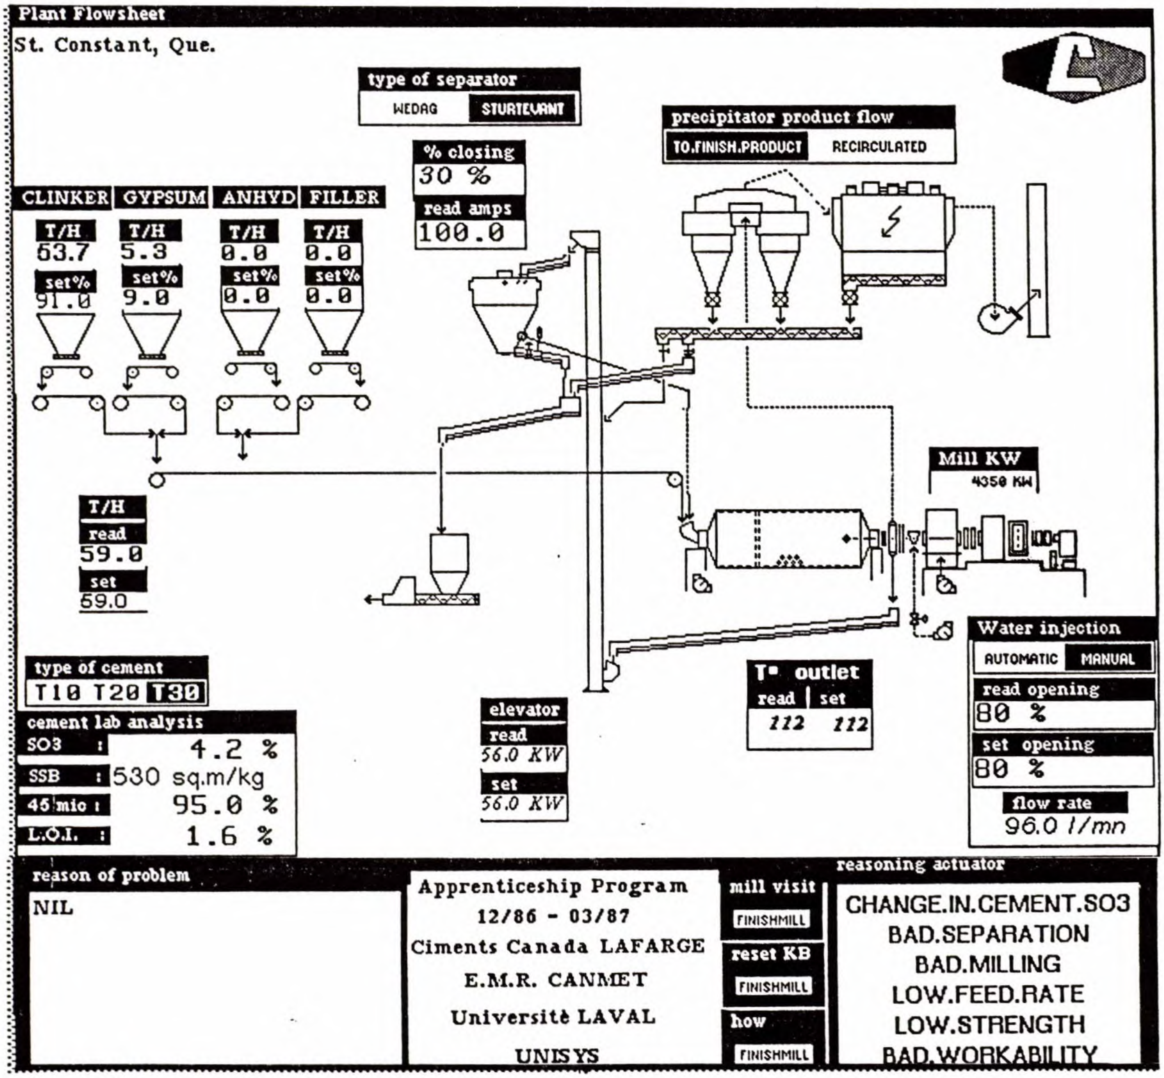

Fig. 1 General user interface of the cement grinding system

extra rules in a negative form. KEE offers also the option of accessing specific methods that are coded in LISP. These methods may perform specific computations or manage to screen or do just about anything that can be done in classic programming except that they have to be written in LISP language. KEE is also a very powerful inference engine that permits forward chaining, backward chaining, as well as the combination of both. The user interface in KEE is very comfortable, especially the fact that the user has access to a collection of active images for variable display as well as to enter new values by clicking the mouse. Figure 1 is a general view of the screen images that are accessible to the user in order to interact with the system. The clinker grinding plant is a conventional part of the process in most cement factories. Its main two components are the grinding mill and a separator. These two large pieces of equipment are also the main sources of problems in the manufacturing of cement. The solids that are fed to the mill are very coarse, their average size is about two centimetres while the final product has an average size smaller than 45 micrometers.

The prototype has been developed to address two big families of problems, problems of production first, which we restricted to low throughput in the plant, and problems of quality in which we included bad workability, chemical imbalance, and bad strength for the cement produced. The prototype functions as follows. The user must first display on the plant diagram of Fig. 1, all the data that is available in the control room. He must also enter the routine analysis provided by the chemistry lab and finally the user must indicate what type of problem has been observed. First, the system is going to make use of the data available in the control room, and attempt to diagnose the cause of the detected problem from this control room data. If the data available to the system is not sufficient, the system will attempt to organize a reasoning in a second phase. As you know such a plant is regularly checked and some plant reports are prepared whenever the equipment is inspected. These reports are of prime importance for someone who is troubleshooting and therefore it is essential that these reports be also accessible to the expert system. The expert system will attempt to diagnose the nature of the actual problem from observations made during the last physical inspections of the plant. In the third phase, the system may require some extra information from the user in order to complete its reasoning. This extra information requested may consist in extra chemical assays requested from the lab, in a visual inspection of some parts of the plant, in a complete shutdown of the plant for visual examination of the inside of some units. If the system fails to reach a conclusion it undertakes a thorough cross-checking of all data supplied to try to detect data inconsistencies. Our current rule system includes some 250 rules that are used in a backward chaining mechanism for process diagnostic. It is also possible to use the same system in a prediction mode by forward chaining through the rule base. We have tested this application of the prototype for a few cases and the results were promising. One limitation however has been observed, it is that in some cases it is required to add extra rules in order to allow the chaining to proceed in a forward or in a backward fashion through the same rule base. This opens the possibility that an ideal solution may actually consist in developing two separate rule bases, one for backward chaining or process diagnostic and the other one for forward chaining or prediction. We have also tested the implementation of side computations in LISP methods and although we found this procedure to be rather cumbersome, it is nevertheless available. Finally, we studied the possibility of differentiating between different plants using the same process but different units such as separators and different sets of rules attached to the particular type of separator in each plant.

My time is running out and I am going to have to conclude on this presentation. First, we can conclude that the potential of expert systems in the cement industry are very important and diverse. The type of application we can consider goes from process audit to automatic process control, however, the display of such a new technology requires substantial investments in persons and in dollars from a company and we believe that a very serious assessment of the implication of an involvement in such technology needs to be undertaken by upper management in the group. We have fully appreciated the power of the development tool such as KEE but some drawbacks that have been mentioned above should also be remembered. As far as the hardware is concerned, our observation is that the two megabyte active memory that was available on our explorer was too small to accommodate comfortably KEE and our particular database. I must also underline the speed of the development since in close collaboration with Frédéric Flament of Laval University, we have defined and entered some 250 rules in a period of three months. Thank you for your attention.

EXPERT SYSTEMS IN CONCENTRATORS: ECONOMIC INCENTIVES

G.M. SWINKELS,
CONSULTING ENGINEER,
VANCOUVER

I want to look at the economic incentives of the application of AI technology in our industry. I elected to take my examples from concentrators, that is mineral processing plants, but my remarks will apply equally well to other types of operations such as roasters and furnaces. They will become less and less applicable to other industries, such as the chemical industry, that have altogether different kinds of problems. I am restricting myself to process control. I am not going to talk about problems related to maintenance or business or financial or exploration. As a side remark, I would like to say that artificial intelligence is a terribly vague expression, and as explained in an earlier presentation, I would rather use the word symbolic computation, since we are talking about computing with words as opposed to computing with numbers. For someone working in operation, the word "pattern" is also very important since one of the most important problems is to recognize the occurrences of certain patterns. This is why artificial intelligence is so important, its nothing more than reproducing what the operator uses in his everyday work. I don't want to repeat some of the key concepts that Daniel has explained in his presentation but I would like to put the emphasis on certain concepts. One of the important qualities required from a expert is that he knows his limits or in other words, his performance degrades gracefully. He knows when to tell you I do not know, please go see somebody else. Therefore, an important property of expert systems will also be to degrade gracefully.

Among the various tasks that expert systems can perform I want to emphasize diagnosis. Diagnosis is a word that can be understood in two different ways, in a medical diagnosis, you feel sick and you go to a doctor. The doctor has a very important piece of information, it is that you showed up to show you were sick. In process diagnosis the situation is somewhat different since it is as if the doctor was walking around and observing symptoms and declaring to the patient that he has to see someone in the hospital because he shows the symptoms of being sick. Process diagnosis is something expert systems are very good at as well as in heuristic control. You have heard this term before, and I will repeat here the definition of heuristic control, it is a control that is not based on numerical algorithms and mathematical models but on rules of thumb. Heuristic control uses concepts such as a little, a lot, very much, which are also called fuzzy concepts. Another task that I wanted to point out is tutoring and assistance. The computer and the expert system can be a teacher or an assistant. Finally, I want to mention a few proper applications of expert systems for instance, if you don't have enough skilled people, enough operators, and another very common instance is when you need organizational memory (Fig. 1). You have all this expertise gathered in your organization and then somebody quits, and you're going to lose that expertise. How do you functionally try to keep that expertise. You put it in reports which very few people read wherever they are. This is why an expert system is a very good place to store that expertise.

PROPER APPLICATIONS

- **not enough skilled people**
- **too many possibilities**
- **need for combined expertise**
- **need for excellence**
- **need for organizational memory**

Fig. 1 Proper applications of expert systems

SOCIAL AND GEOGRAPHIC CONSIDERATIONS

- few plants in isolation
- small staffs
- shift work
- turnover

Fig. 2 Social and Geographic considerations to apply expert systems

At this stage, I would like to submit a proposition. Control is a complex operation, and a good operator is an expert. From that proposition I want to show you that the way we control our plant presently is not proper, and the use of expert systems could improve process control. There are some social and geographic considerations that explain why it is difficult for a mineral processing operation to keep its expertise and have it available at the proper time and at the proper place (Fig. 2). Most of our mineral processing plants are isolated contrary to chemical or petrochemical plants. The market for expert operators is therefore very restricted in our industry. Isolation may have two different meanings, one could be that we do not have a similar type of operation nearby or that the plant is physically located thousands of miles away from urban centers and it is even difficult to bring operators to the operation site. Also economic realities force concentrators to operate with small staff which means that it is very difficult to have all the expertise on site. The fact that operations are run on a shift schedule makes it also desirable to uniformize expertise between shifts. The mineral industry is also experiencing a very high turnover of manpower, very often due to the very difficult and remote conditions of the work. Every time a new operator is coming on staff, the company has to invest in training and this investment is lost when the operator leaves. This is another reason why expert systems are a very attractive answer to these social and geographic considerations. We now move to the technical arguments.

The first point I want to make is that the success of process control has not been as good in mineral processing as it has been in refineries. One of the chief reasons for that is that feed materials that have been provided to us by mother nature are of very different compositions. What is even worse is that very often we cannot even measure and categorize the variations of this feed material. We can therefore make constructive critique of classical control in order to show some deficiencies that expert systems could remedy.

The question is not to use expert systems instead of process control, but to use them jointly. If we look at what process control does, we have a black box there that solves a second order differential equation over and over every 40 seconds or every 10 seconds, depending on what you want, in order to effect changes to the process. As you know, you have also those instrumentation engineers running around the circuit and tuning these controllers in order to take care of certain variations in the feed material. The parameters against which the controller are tuned are either the tonnage of the feed or the variation in the size of the feed or some other property of the material in the feed, but it is impossible to tune the controllers to take into effect all the possible variations in that feed material. This is where the expert system can do something. Some changes in the operation can be taken care of but if we look at a case for instance, where recovery is a second order function of say the quantity of frother or of cyanide used as depressant, if that curve shifts then we have a very nasty situation. The controller may assume that it is working on the rising side of the curve and assume that if I increase the quantity of cyanide then the recovery is going to come up when, actually, due to a shift in the curve, it is working on the declining side of the curve, and by increasing the cyanide, it further reduces recovery. The process control assumes a model which in most cases is not true in reality. If we look at the very conventional process control systems, they usually have all in common one feature: they do not degrade gracefully. At one point, the system is acting in a way which has nothing to do with reality and the operator has no choice but turning it off. One mode of operation that is often used is that people assume a relationship between grade and recovery and they aim at

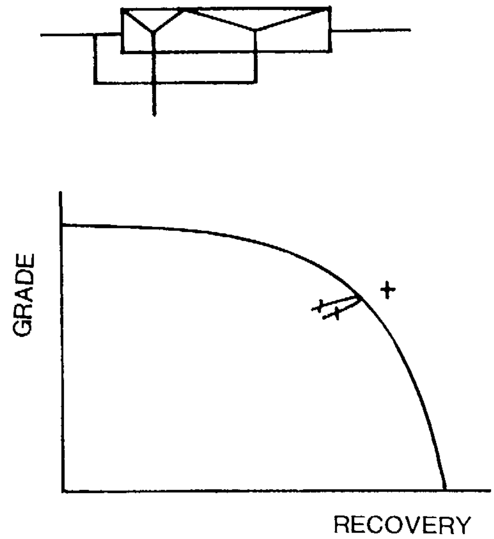

Fig. 3   Conventional grade - recovery curve

a certain grade and a certain recovery (Fig. 3). Choosing a target point outside this grade-recovery curve may lead to considerable problems in conventional process control. Another dramatic situation occurs when you have a process breakdown, in that case conventional process controls leads to further complications. Of course, process control systems have a set of alarms to face such situations, but these alarms are geared towards only a few possible breakdowns. This is why an expert operator becomes useful because once he has detected an alarm, 10 minutes ago, he is expecting something else to occur as a result of this alarm. If it doesn't occur, he is going to suspect that something is wrong in his system. This type of reasoning of course cannot be done by conventional process control. Neither can it be done by an operator who is having his lunch break. Another limit set on the process control system is the lack of sensors presenting the sufficient ruggedness to operate in our environment. The process control system cannot be useful during the breakdown periods or during emergency situations. There are therefore a number of areas in which conventional control systems do not provide assistance to the operator and this is where expert systems can play an important role.

We finally get to the point where we can describe what expert systems can do better than conventional control (Fig. 4). One of the first contributions of an expert system could be to check on the consistency of your conventional process control. If you know that taking a certain action you expect a certain result, the expert system can verify that this result is actually occur-

EXPERT SYSTEM CONTROL

- **check on consistency**
- **check on control**
- **actual control**

- **diagnostic component**
- **procedural component**

Fig. 4 Applications of expert systems to process control

ring and, if it is not, then interrupt the action of the conventional process control. Another role of the expert system is to check on the control system itself to verify if the process is still operating in the mode assumed by the control system. Another contribution of the expert system may be to do the actual control of the process. Such applications are already reported in the control of kilns and in the control of cement plants. I would like to mention two more components of potential applications of expert systems. One is a procedural component where the expert system assists in complex procedures, such as the procedure for starting up a plant, and the second component is a diagnostic component which is more powerful than the strict application of rules. The diagnostic component I am referring to here is one that maintains a concept of what the plant should look like. This is a very interesting aspect because this is precisely where the kind of work that has been done by CANMET in the SPOC project would come back in its full importance. There are expert systems that are operating in real time in plants that belong to both EXXON and SHELL and probably to others, although it is difficult to know because this is not the type of information that companies share easily.

Now I would like to give some consideration to the way an operator on the floor interacts with an expert system. At the present time, with conventional control systems, there is no interaction. The control system is based on mathematical models with which the operator has no interaction. He sits there and lets the system do its thing. With an expert system, the operator can ask for an explanation and get an explanation. Also rules that are discovered by the operators, can be later added to the system and the system then would behave according to rules they are familiar with. I suggest that the incentives for expert systems development have to be found first, and foremost in

GENERIC TASKS

- signal understanding
- diagnosis
- information fusion
- design and planning
- heuristic control
- intelligent agent

- tutor and assist

Fig. 5 Generic tasks that can be assigned to expert systems

preventing errors from taking place. The training of operators is another important aspect (Fig. 5). There is an enormous amount of money being spent in either training operators or paying for the mistakes resulting from a lack of training of operators.

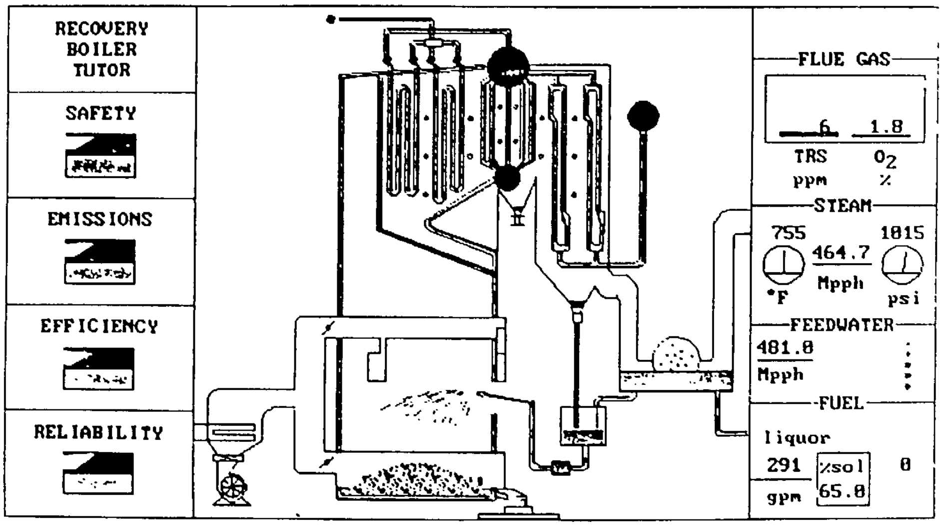

Fig. 6 Example of user interface with expert systems for the control of a boiler

I would just like to conclude by presentation by giving the example of an expert system that simulates the control of a recovery boiler in the paper industry (Fig. 6). This was a project commissioned by an American pulp and paper organization and this expert system is used as a tutor to train operators before they actually have access to this very important and expensive piece of equipment. Thank you for your attention.

PRESENTATION OF THE AFTERNOON SESSION

D. LAGUITTON,
CANMET,
OTTAWA

Good afternoon Ladies and Gentlemen and welcome to our afternoon session. We are going to listen this afternoon to three presentations by suppliers of AI. They are the people who are promising us the results from applying their technologies and I will invite the users to take this opportunity to make sure they ask all the questions they have on their mind related to their field. I would like to read a few quotations that may bring some comments or debate from the afternoon speakers. One is a quotation from Datamation, March 1, 1987, it is extracted from a report on an interview given by Michael Blumenthal, the new chairman of UNISYS. The author of the report writes, and I quote "the company has shelved for the time being, artificial intelligence-based products plans that had originated at both Burroughs and Sperry". This is an important statement and since we are fortunate enough to have a representative of UNISYS, we will ask him to comment on this quotation. Another quotation, which I don't have in front of me, but which I remember, is extracted from Automation News, February 16, 1987, in an article questioning whether AI was really being transferred from the Lab to the plant. The author discussed the fact that several AI companies have shown some plateau in the profit curve in early 1987. The quotation goes "would AI be like a porpoise, stranded on the sand bar, intelligent but unable to go anywhere?". I would invite also comments from the afternoon speakers on this matter. Finally I will read a quotation from Canadian Research, March 1987. Its an article on AI, and you might not know it if I read only the quotation. Here is the quotation. "The conversation you could have with your microwave could be at the higher level than the one you could have with your friend". I regret that this is the kind of mythology that is being pushed around and in that particular case, this quotation is probably more a reflection on the quality of the conversation the author has with his friend than on the quality of his microwave oven.

TECHNOLOGY TRANSFER IN KNOWLEDGE BASED SYSTEMS

D. BROADHURST,
UNISYS CORP.,
WINNIPEG

The Datamation article was based on research that the writer had done four month's previous to the publication and he did not go back over the comments that Mr. Blumenthal had made. In fact, Blumenthal himself never said those things and if anyone wants an official letter signed by our company with respect to that, it can be arranged. However, it does point out a serious concern with respect to "is this is in industry or is this not", and, it is very difficult to qualify and when you look at the facts, this is an industry that's in trouble and is an industry that is going out of business in a hurry. When you look at the opportunities, this is definitely an industry, this is the future of the entire industry. The evolution from data to knowledge processing means we move from 10% of the information which is data, to the other 90%, so our marketplace has just multiplied itself ten times, so in the future it sounds it will be there. There is enough wisdom in the company to understand that our company made a very significant strategic decision to get into this technology. There is enough wisdom in the company to take an assessment point of view and one of the interesting things, in the Canadian program which we have underway here, is making a major impact on that assessment, because of the advances that are happening in Canada at a rapid rate. Canada is actually in a unique position globally to take advantage of this technology. I should qualify also my ability to stand up here and speak to the people in this room. First of all, you know that Unisys is the merger of Sperry, Univac and Burroughs. Those who are on the Univac side, which our Japanese subsidiary still calls itself, say that Unisys means that Univac is still our supplier. We are a large company, 125,000 employees, over 10 billion dollars a year in sales and the reason the merger is important is because we are in an industry where you can be a 5 billion dollar company and you can be out of business in a very short period of time. So the strategy is that the only way to survive in our industry, is first of all through size, secondly that size will allow you the necessary structure and funding to assess technology and to grow. Profit from knowledge: I think this is a key word, and that is why I put it up (Fig. 1). The more highly technical a produce or a process the more easily it can be made obsolete by the emergency of new technology. I spent about 18 years of my life in mineral processing, before I joined Unisys. I was the expert in the field of flotation technology and I thought I would mention this because I feel amongst peers. I spent so much time designing process circuits, setting up factories and doing those kinds of things. The points that were brought-up with respect to the front end of the

PROFIT FROM KNOWLEDGE

The more highly technical a product or process, the more easily it can be made obsolete by the emergence of an entirely new technology or by a more ingenious application of an existing technology.

Fig. 1 Profit from knowledge

product that we get from nature and the industry, and the effect that has on the process technique were very dear to me because I worked in almost any kind of process you could imagine, here in Canada in a small company, installing those technologies in 65 companies around the world. I knew what it was like to go in as a young guy and pick up on corporate expertise, because the senior technician in the firm that preceeded me, had 30 years doing what he had done and I walked in, a 22 years old, and Reginald Schmidt was the senior technician. I went through the knowledge extraction process, I was trying to get out of Reginald Schmidt what he knew about our unique process technology. I also went through the young whipper-snapper attitudes of Reginald Schmidt, etc, etc. Reggie, and I became best friends after about three years but the first 2-1/2 were really tough. That was my first experience with the knowledge engineering process although I knew nothing about AI at the time. You're all reading this: the EnRoute magazine current edition has an article on AI (Fig. 2).

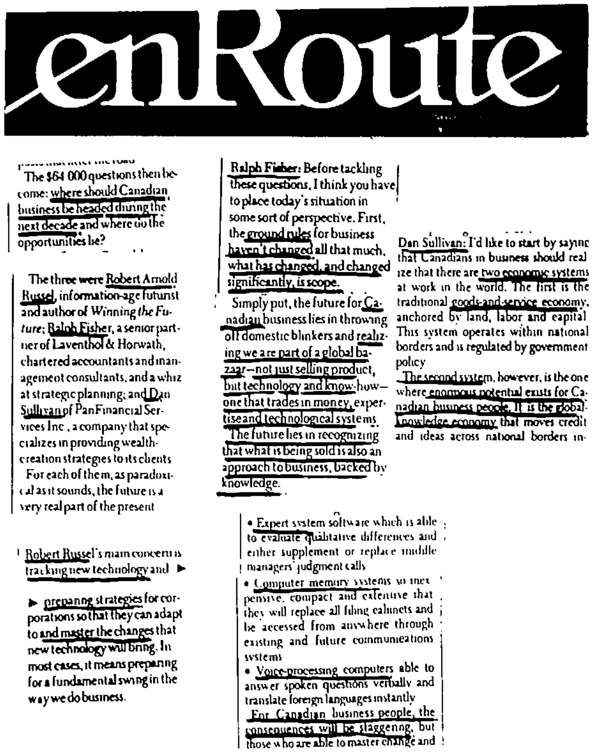

enRoute

The $64 000 questions then become: where should Canadian business be headed during the next decade and where do the opportunities lie?

The three were Robert Arnold Russel, information-age futurist and author of *Winning the Future*; Ralph Fisher, a senior partner of Laventhol & Horwath, chartered accountants and management consultants, and a whiz at strategic planning; and Dan Sullivan of PanFinancial Services Inc., a company that specializes in providing wealth-creation strategies to its clients.

For each of them, as paradoxical as it sounds, the future is a very real part of the present.

Robert Russel's main concern is tracking new technology and ►

► preparing strategies for corporations so that they can adapt to and master the changes that new technology will bring. In most cases, it means preparing for a fundamental swing in the way we do business.

Ralph Fisher: Before tackling these questions, I think you have to place today's situation in some sort of perspective. First, the ground rules for business haven't changed all that much, what has changed, and changed significantly, is scope.

Simply put, the future for Canadian business lies in throwing off domestic blinkers and realizing we are part of a global bazaar—not just selling product, but technology and know-how—one that trades in money, expertise and technological systems.

The future lies in recognizing that what is being sold is also an approach to business, backed by knowledge.

Dan Sullivan: I'd like to start by saying that Canadians in business should realize that there are two economic systems at work in the world. The first is the traditional goods-and-service economy, anchored by land, labor and capital. This system operates within national borders and is regulated by government policy.

The second system, however, is the one where enormous potential exists for Canadian business people. It is the global knowledge economy that moves credit and ideas across national borders in-

- Expert system software which is able to evaluate qualitative differences and either supplement or replace middle managers' judgment calls
- Computer memory systems so inexpensive, compact and extensive that they will replace all filing cabinets and be accessed from anywhere through existing and future communications systems
- Voice-processing computers able to answer spoken questions verbally and translate foreign languages instantly

For Canadian business people, the consequences will be staggering, but those who are able to master change and

Fig. 2 Reproduced from EnRoute Magazine

This is an important one. We've got some people here like Russel and Fisher and Sullivan. What they are indicating in these articles is that the key thing for people to do in Canada is to track technologies that are going to create opportunities in this nation in the evolving local economy. We know there has been an article in the Globe and Mail in the last few days. Alcan is making statements that the metals industry is weak, you know mining is a stable industry but you know we look at the real statistics and see that it is a question of how much can we depend on our resources and we know that we are moving towards a goods and service and a knowledge economy. The key thing is that the expert systems technologies are the players that are going to allow us to take advantage of that economy or be a player in the economy. If we use them today they will give us a strategic opportunity in this revolution. So our analysis of the situation is not whether AI is over promoted, we don't really think it is, even if there is a seminar every day and a half in Europe and probably two seminars a day in the states, but rather that it is under deployed (Fig. 3). We don't see people using the technology. The examples we have here today are really significant because they are starting, we are seeing the first movements into your industry. Maybe next year we'll have the figures return on investment of applied expert systems at CEZ and

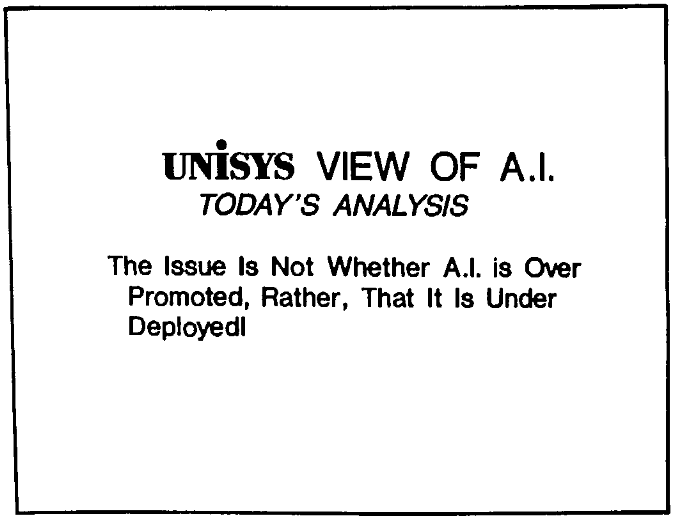

Fig. 3 Unisys view of A.I.

Lafarge. We have, you know, the return on investment for the system we built with Northwest-Orient Airlines, wherein they use an expert system to price economy seats. If they can raise the level of one economy seat to full fare for 25 per cent of their flights, the potential is for a 12 million dollar increase on the bottom line. In the airline industry, there is no potential for additional passengers, there is no potential to cut the cost to the fleet operation. There are no areas where you can improve what you are doing or very, very marginal, so you have got all this merger fight bankruptcy scenario going on because its over supplied. So how do you then provide performance in your corporation so that you can grow. The only place is efficiency, and the one area that Northwest has applied it on, was with respect to seat pricing. Our company made the first connection of an inference engine to an 1100 main frame computer. Now all the other airlines are giving in to that wavelength. To illustrate how difficult it is to conduct technology transfer, the first thing is, that, to really begin to understand, to teach somebody how to tie a knot takes about a month. Now, when you are a kid and you learn how to tie your runners takes about a month, and that is because language is so limiting. So technology transfer is a very difficult thing to make happen. One thing you are learning today is that expert systems is a technology not a product and I would guess that a lot of you this morning did not know that when you came into the room. You thought it was something you could buy off the shelf. What we've got here is a technology, it is pervasive, it can work in medicine, it can work in mining, it can work in finance, it can work in all kinds of areas. We found out this was a technology, when we went to 3M Corporation when we decided to get into AI and they told us that yes to them it was important, it was more important perhaps than chemistry, and they have 38,000 chemists and they said "how do we build disposable expert systems"?. 3M makes bandage, scotch tape, and sand paper and all kinds of things, so they know what "disposable" means. We have technology that could solve huge problems, like the NASA space shuttle problem with their scrubber, the oxygen scrubber, that they couldn't solve in eight years, and they solved in three and a half weeks, using AI. We had a technology that was disposable and we found out that you can solve a problem, using this technology, and the problem can be something that you don't have to solve again.

To review briefly the skills of people who use expert systems. I put management in there, and the reason I put that in is because in the beginning we were looking at AI, first of all, as a computer company and we didn't understand its power as a management tool. Top management decision making tools and advisory tools are really early forms of AI. Operations are what the focus has been today and probably the other areas as well. The operations

The Elements of an Expert System

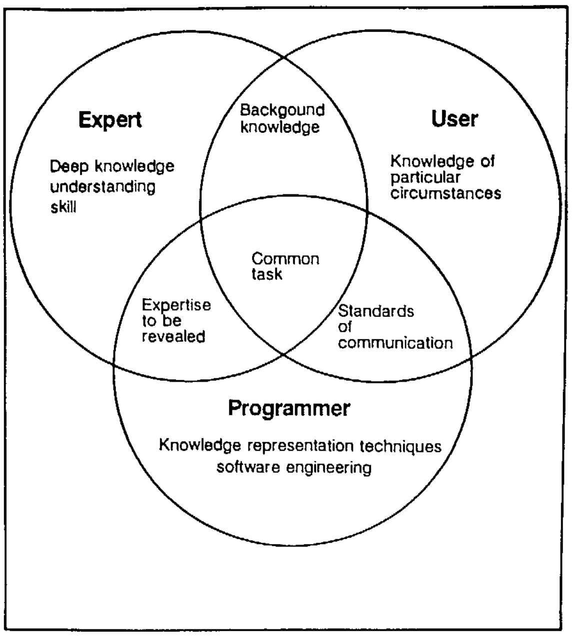

Fig. 4 The elements of an Expert System

are where the early things are starting to happen. You know Unisys employs a lot of good engineers. If you look at the structure of our company, you find out that most of our people come from Detroit and Minneapolis, based heavily in the engineering community. What we are doing is we are engineering the use of this technology, that is what engineers do, they take a technology and put it to use. Our dedication is to see this technology in use, we are not interested in advancing this technology. The universities, the MITs, the Carnegie-Mellon's, the Stanford's the Waterloos, the U of Ts, etc, are advancing the technology. We are interested in taking that advancement, and putting it into practical application.

Three priority elements are relatively obvious by now (Fig. 4), the user is one we haven't talked about too much, it is one I will focus on, because thats the back side of the technology transfer. There are two sides to this, two directions, one is transferring the technology in so that somebody understands it, then the second side and a more important on a long term is how does the user use the technology that is there. Transferring the knowledge in a knowledge system back out to the user. In the next figure (Fig. 5), there is a line that says that in 1984, there were about 60 PhD graduates with actual AI development experience, this is a Transport Canada study, so it was evident in '84 and I'd say in '85 it was doubled, which was 120, lets say in '86 we got up to 300 and in '87 we got 500 university graduates in all of North America, trained in AI. We need probably to get Canada to use this technology, 10,000

Transport Canada: Expert Systems — **The Supply of Knowlege Engineers**

Industry

One of the trends developing in the industry today is the determination to generate AI experts, in particular "expert systems engineers", through in-house projects. This is based on the recognition that there will never be experienced AI researchers, let alone AI project leaders capable of undertaking a major expert system project, available from the universities. It was understood, as of mid-1984, that there were at most 60 Ph.D graduates with actual AI system development experience leaving university in the entire North American continent annually. At the same time, it was estimated that the annual requirement for AI researchers and engineers of the calibre required for actual system development was greater than 300 in the U.S. alone.

Fig. 5 The supply of knowledge engineers

people in two or three years. Our corporate target is to have 10% of our technical people in AI by 1990, so we need a lot of people, so it is evident to us that academia couldn't supply the requirements, the human requirements, to support developing systems. So we took the corporate position that perhaps we could assist in that process and we really set up in the States, what we call a knowledge system centre. That centre was supposed to understand the technology and use the technology to get our people up and running but that centre is now evolved to a centre which interacts with the customers. Basically in the centre there are trained manpower, a scaffold of equipment to work on which is a key thing, and we bring projects in from the customer base. When I looked at Canada, and was handed the mandate to start a Canadian program, I said well in this country, if we do this in simply Toronto, or

Montreal or Vancouver, isn't going to suffice. The geography hit me and I thought, well living in Winnipeg, I understand the regional disparity to the finest degree, I would set up a network of centres in Canada. So we started and we've come a long way. Now in Winnipeg we are opening a centre in partnership with NRC. We'll have five machines in this centre, we have a centre in Mississauga already established which is our corporate centre, we've created centres with universities in Ottawa here, at Carleton, and in Halifax with Dalhousie, in Hamilton at McMasters, and recently, we are currently in the process of putting in a centre in Victoria and in Vancouver. The interesting thing is that whether we call them centres or units, they are all dedicated to the same concept. The concept is not to develop expert systems, the concept is not to further the technology and do research, but rather to transfer technology to people who want to use this. So the big mushroom is how do we get started in expert systems. That is a paramount question, that is like how do we market in China. It is the only parallel that I could think of, that would be as difficult to comprehend. It takes a long time, there's a lot of technology, etc., etc., so what we devised is with the best way to do this is by walking people hand in hand through a project, as we do here with Laval and with Lafarge. In the apprenticeship program, we have formalized an approach to getting people involved in AI and typically there are three people in the program, the knowledge engineer from our corporation, the person in the customer's company to apprentice and the other person would be developing expertise for the user in this case (Fig. 6). Prerequisites to participate in the program is that obviously there has to be tool training

The Apprenticeship

The apprenticeship provides twenty days of Knowledge Engineering consulting and support from Unisys over a period of three months. This consulting and support are generally provided at a Unisys location. These twenty days are divided into several week long sessions in one of two schedules.

Fig. 6 The apprenticeship program

for the people involved, so they have to get their hands on the equipment. We have formal industrial education in these tools, which we offer in the program. And the first step obviously is to set up a meeting to go over and look at the potential domain areas that the customer would like to get involved in. We do this generally in a section with multiple knowledge engineers that have a broad base of experience. We have a about 80 internal projects underway in the corporation and I'm not sure what our count is, but we must be close to 30 or 25 in that range of apprenticeship programs now that either are about finished or in process of being conducted. So it is a broad base of experience, as to what type of domain areas we are addressing. There are two types of schedules that we go through, it depends on the type of problem, the availability of the people and that sort of thing. Generally, it is in the knowledge centre and then at home type of process. In the knowledge centre

Apprenticeship Program Results

At the end of the apprenticeship program the apprentice has built a working prototype expert solution to the selected problem and management has enough of a result to evaluate its value and potential benefits. The apprentice is also trained in the application of the KEE and Explorer environment to solving problems in the customer's business area. These results are accomplished in three months with the help of the apprenticeship program rather than the year or more that it would take unassisted.

Fig. 7 Apprenticeship program results

you're in an environment where there is a lot of people with good expertise available to you, at home you work on your own and you bring the results back and that sort of thing. Conventionally what we find out is that the people never bring enough information to the first session. Their thinking is more along the terms of conventional computing and we find out that after a day we've run out of basic information, or a day and a half, because we can code so many things so quickly and once you map out the direction you can really get going. At the end of this program then we have built a working prototype solution (Fig. 7), the prototype is not a finished project but one with enough meat in it that management can make an economic decision as to whether to continue with the project or drop it or whatever. Okay, we've trained the people in KEE/Explorer environment and we've done this in a reasonable time frame. In one business quarter, we've got the corporation that far. Now often by the time they are there, they don't even appreciate what's happened in one quarter, but, a lot of companies don't accomplish anything in a quarter. So we think its pretty significant to take a corporation and give them a full hands-on understanding of the technology in a working prototype in one quarter.

This area is one which I think is extremely important, and for those of you, I put this together because I felt most of you were primarily in engineering areas, and as I suffered along in that area for about 20 years against the accountants, I thought this was relevant (Fig. 8). Management's role in cost justification of new technologies is to adjudicate between the engineers and

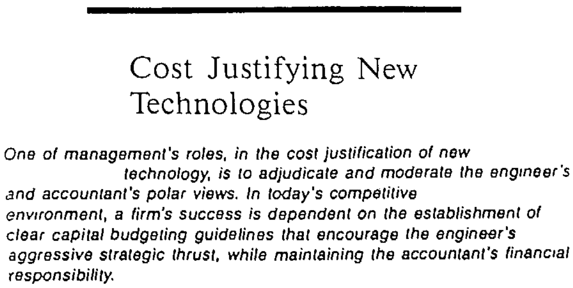

Fig. 8 Cost justification of new technologies

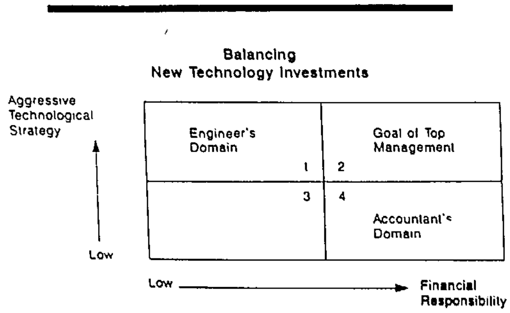

Fig. 9 Balancing new technology investments

the accountants point of view. But the firm's success is dependent on establishing budgetting guidelines to allow a strategic thrust to the company, and maintain the accountant's financial responsibility. The following diagram (Fig. 9), shows you the accountant's domain bottom line, high on financial responsibility, top left the engineer always wants to do some wild and wacky new thing, and there's management looking at these two guys. Right now I would say Canadian management is sitting near section 2, kind of more toward financial responsibility, than he is, well not really at either extreme. We want to move him now into an equal position. The books are secure in an equal position but the whole company is at risk because it is not aggressively addressing technology. This is a common problem, that is very difficult with so many things that are involved in this. Basically what we're saying is you might get knocked off tomorrow by an organization that you don't even know exists today because of technology. And if you're not tracking and keeping track of what's going on, and have a budgetting process that will allow you to address new technology in your company, you could be in trouble. The next figure (Fig. 10) was produced by Western University, in a study on actually

Manufacturing Automation: Promoters and Barriers

Promoters	Barriers
• Senior Management's Attitude (65%)	• Total Costs Involved (63%) • Uncertainty re Benefits (54%) • Availability of Capital (46%) • Expertise Required (47%)
• Rate of Change in Technology (28%)*	• Perceived Payback (45%) • Union or Employee Group Relations (32%)
• Vendors Reputation (27%)*	• Employee Attitudes (35%) • Technical Requirements (38%)

*Marginal

Fig. 10 Manufacturing automation: Promoters and barriers

on manufacturing automation but they really, I think, are symbolic of the type of university studies that indicate to us how companies address automation in high technology. The promoters on this side, senior managements attitude is the key promotional attitude. If senior management is technically aggressive, that that is the direction the corporation will go. Or if a government department is technically aggressive, that's the way they will go. So you go into one government department and see office automation from top to bottom, and you go in the another government department you see paperwork top to bottom. Note that marginal is rate of change in technology and the company's reputation, again this is on manufacturing automation. The barriers are obvious, costs, uncertainty of benefits is the big thing, I think that's the biggest barrier to penetration of AI. Where is the benefit that will really come down, because we have to deliver to you a lecture on all the marvellous returns on investment. The reason we use marvellous there because the guys that are making money using this technology are in a strategic winning position, they're not telling anyone about it and there are quite a few of them, over 150 at least in North America. The following series of figures is on managing technology transfer (Fig. 11,12,13,14). I presume we have an expert system, all we have to do is to transfer the technology out of that.

Manage the technology transfer

Plan at the outset how you are going to transfer the system to users.

First, get the users involved early. A main reason to encourage end user development of knowledge systems is to side-step these transfer problems. A person who develops his own system is more likely to want it and use it.

Fig.11 Management of technology transfer

Manage the technology transfer

Second, deal with people's fears. Some people are threatened by a system that can make decisions. They see it as a threat to their job, not as a tool to help them. People tend to be more fearful of expert systems than traditional systems.

Fig.12 Management of technology transfer

Manage the technology transfer

Third, be wary of developing systems that will disrupt an organization or cross organizational boundaries. These are much more difficult to introduce successfully.

Fig. 13 Management of technology transfer

Manage the technology transfer

Systems fail due to indifference on the part of the users and irrelevance to the business. Either a vice president or an officer of the firm must *really* want the answer to the problem being solved, in order to get the system into use, once it has been developed.

Fig. 14 Management of technology transfer

One of the key things is to get the users involved. There's two basic strategies to getting into AI and I could use corporate names but I won't. One large aerospace company uses the strategy of taking people, producing AI gurus, and put them through a whole PhD process in AI, so they take the philosophy and all the dynamics in a real deep structure. That takes them about a year, a year and a half, then they come back to the organizational structure, they've been out of the company for a year and a half, and when they come back they are no longer that relevant to the line operation that they came from. The success rate in that particular corporation of applying AI is around 8%. In our organization, we found out that the best tools that we had to offer as a company, are the tools we can develop ourselves in-house, the tools that our own organization actually built, so we identified the user is the key. In computing we talk about things like language generation, fourth generation, third generation, fifth generation, fourth generation languages are end user tools and they are really accepted well because the end user's in charge of his own project, so it is keeping into account the line of operations involved. That's where the apprenticeship program comes in, because you take the line operation, you know, we take the circuit grinding guy who the expert in that area, he's the line operation guy. He's involved in building the tool, if that tool is put into the plant, into use, he has ownership, pride, all those key essential human elements to make that thing succeed. It's sort of assuming in our approach strategically, that if you are going to get into this technology, you should design your approach for success, don't design it for failure. Second, deal with people's fears, like Reginald Schmidt, he was really afraid of me even though he was retiring. A lot of people are threatened by a system particularly one that makes decisions, I mean that's getting a bit scary. They see it as a threat to their job, that is true, they are more fearful of these systems than of conventional systems. I think what you have to do is really have a top prospective, now we're talking here about putting systems into production, we're talking here about applying this technology in reality, I'm not talking about building research projects here, these are real human issues, and this is going on. So you take the loan officer system thats going into American banking right now, a loan officer has a stack of rules that he has to live with. Now, there's two things, a) if you have a lot of loan officers, say 4 or 5 thousand of them, you have to bring them in every six months, to go through a two weeks update to keep them current with whats going on in the bank's policy. That costs the bank a lot of money, b) the guy may not be that good, c) if you want to change policy on loans, you can't do it instantaneously, its another six months before you can bring the crew in and reeducate them. When a tool was developed for that situation, an expert system, with all the rules governing loans in a particular bank, an advisory system for the officer, the thing that occurred was, that a) the travel was eliminated, b) a new diskette could go out next month, c) the loan officer then had an expert advisory system, he could come closer to the customer than he could before. What happens to your loan officer as soon as you get too close, you know what happens to the friendly bank manager, you know what happens to him, he moves. As soon as you get to work with the guy, as soon as he understands whatever you're doing in life, in your business, in your relationships and whether your good or bad, they move him to another location because, he might give out a bad loan because he's developing this cordial relationship with the community. So the bank spends all its money impressing you now with their image where they used to impress you with their people. This technology is offering the bank the opportunity to go back and impress you with their people and put the dirty work on the machine. So for the banking industry, they see real significant evolution in the way they approach dealing with people. I was very pleased to see the deputy minister here this morning, that type of direction is an inspiration to people doing very difficult things.

REAL TIME EXPERT SYSTEMS FOR THE PROCESSING INDUSTRIES

J. DREUSSI,
L.M.I.,
CAMBRIDGE, MASS.

If you think about the expert systems technology, Mr. Broadhurst is talking about, it has a common thread throughout, two concepts actually. The first thing is every application that he talked about is real time. Every expert system that is going to be developed in the future, is going to be real-time. The expert systems we think about today are interfaced with human beings sitting at a terminal typing in data. That data goes into the machine, is good forever. In real time application thats not necessarily true. Data is only good for a certain amount of time. The other thing that he mentioned was the importance of given the end user of the expert-system shell. What I'm going to talk about today is real time systems, what it is going to take to, or what it takes to implement real time expert systems, and what we would like to see on the user end. Maybe we can do away with programmer and the knowledge engineer and just have the expert and user sitting down in front of the thing. I'm here to discuss the real time expert system application package that is called PICON. I'm not going to sell PICON, I'm just going to tell you some of the things that are in it and the philosophies behind it. What does it take to have a real time expert system and how is it characterized. First of all the large number of inputs maybe up to 20,000 points, maybe more, data points, sensors, whatever, data management, data base, data bases, whatever it takes to get the job done. Very large knowledge bases, up to 100,000 frames, maybe more, rules, describing rules and objects. Some of the frames will be active all the time, some of the frames are only activated when certain conditions occur, so you focus in on them. Time, what does time have to do with real time systems, sometimes we talk about a real time system as being real time, as being fast enough. Real time is more than just fast enough, real time means I understand what data comes in, how long it is good for, and I understand whether or not to use it. As far as the user is concerned, if you are putting together a system, you have to have a way to retrieve all the attributes, you have to have a way to retrieve the knowledge base and pieces of it based on attributes of your objects, based on attributes of your rules, based on whatever you dream up. You have to have a way to go back and change the thing, one of the characteristics of expert systems is why we have expert systems and why this technology has caught on is because they are flexible. They are flexible, we can change them, we can grow, we can evolve. And finally, if you are going to have a system that is this big and perhaps it's going in a critical area, where either safety or money is involved, you want to be able to test it and simulate it in advance, you don't want to go off and put the thing on line because you never know what the consequences are going to be. In addition to all that, if you are going to have a real time system, you have to be able to interface it with different kinds of control systems, different kinds of main frames computers, you may want to go with some kind of data highway or network or you may want to use direct bus to bus connections with it. Well, if we're going to have an expert system in a factory, consider a factory to be anything that is real time, any process can be a factory, each time we make decisions we go through the same kinds of procedures as a factory goes through to make decisions. When you're in a factory, whenever you're in a system, you're working with a system where you can have 20,000 points you have 100,000 different rules in there, you've got 200,000 or 500,000 objects to worry about, what happens when an alarm goes off what happens when something occurs that is wrong? You made a decision for example to go off and buy some oil at a certain price, how's that going to affect the rest of your operation, what's going to happen with it. Supposing you make the wrong decision, how do you go back and back off from that or if you're working in a power plant and an alarm goes off, which triggers hundreds of other alarms, how do you focus on that, how do back your way out of that. The systems have to be able to handle hundreds of these alarms at once, and not confuse the user. You have to allow the operator to grow this control software in the expert systems, just as they grew the control software in current processes they are working with. We didn't get to semiautomated factories overnight, it took a long process, and that's whats happened with the expert system technology, its a long process, in a real time arena this is very slow because people are very reluctant to let a machine make decisions and take control of the plant. And finally, we need to train our personnel as to why we want to do this. Whenever we're designing an expert system, that is going handle 100,000 points, has thousands of rules and thousands of objects in it, what kind of a computer is it going to take to scan all those points, and to fire off all those rules every time you need to fire it, every five seconds, every five milliseconds, or whatever? In putting one of these real time systems together, what we get is we look at how a human being does it, we look at the human in the control room for example, whenever everything is running fine in a control room, whatever your process is, the operator is sitting there and he's just looking at a few gauges, a few dials, a few parameters in time, just looking at the top level of it. But if something goes wrong with one of those gauges, gets outside of range, what happens? Well at that point, he focuses in on the other part, on another part of the process. For example, if you've a furnace operating something that is temperature or pressure oriented, as long as the temperature or pressure is ok, you don't worry about what happens downstream to the pipes, you don't worry about what happens to perhaps fire hazard downstream from the furnace. Whenever, the pressure gets too high or the temperature gets too high, or vice versa, or they get too low, then you go down and you focus on a particular area of the system. This is a different kind of chaining mechanism, than the one we are used to talk about in an expert system. Lets look at a conventional expert system diagram, this is generally how we compare expert systems, conventional expert systems, to a conventional piece of software. When you look at a conventional of software, you've got to program in all the knowledge of your domain stuffed into one box, then you've got the data interface with that thing. When you look at an expert system, in an expert system you have your program which is called an inference engine, and you have the knowledge about your domain separate from the inference engine, and this program is strictly a data driver. When data comes from the outside world, some rules and heuristics are worked on and then it fires the inference engine. Well, if you look at a real time application, what's it going to take? The first thing we need to recognize, is that we're going to have some engineer experts who are going to built a knowledge base for us. If this thing is going to be real time, then we're going to have some way to look into our process, and finally we're going to have to have a way to talk to our operators up here and give them advice and have them come back and ask why, why did you tell me that, what is the significance of that, and we're going to have to have all kinds of chaining mechanisms inside this, that include forward chaining and backward chaining and a few others that aren't talked about. You take that structure of a real time expert system and break it apart you get this, down here is a real time system that might be running in a conventional computer of some kind, up here the inference engine a system that might be running in some kind of a symbolic processor. What you need, what we need is for this real time processor to be able to call the data highway, the network, pull in the data, stick it in the memory, and do some low level inferencing on it, stick it in the memory, have the inferencing pick that stuff up and combine that with the knowledge base and then send the display up, send the results up to the operator to display. So our basic expert system structure gets to be a little more complex when we start talking about a real time application.

What are the steps involved in developing an expert system. First of all we have to have the knowledge acquisition, we have to get the knowledge from our expert into the machine, then we have to develop rules about that knowledge, and then we have to put the interfaces together. Well, if you look at the knowledge acquisition, if we would like to have a system that the end user can use directly, we would like to have the system represent icons, plant elements or what have you in the end user's language. We would like to have the end user be able to draw up the schematic of his actual process in his own language. A schematic looks something like this. This is a schematic for a simple fault isolation process. We've got a holding tank, and we've got some filters and controllers and we've got a lot of sensors out here, there's F1, B1 and C1, there are various kinds of sensors, and what we would like to user to be able to do is sit down and build up a schematic of his plant and once he's laid out the schematic of his plant there are certain logics in here that he gets values from, those are the sensors. For example the level of his holding tank is a sensor. He goes in, he builds up his schematic in the language that he understands thats his language right there. So what the engineer does, is he selects the icon type, the icon vary from industry to industry, he defines the sensor values, he defines for instance a tag name for his sensor value, which is the name of a particular sensor of the network. How often are you going to have to make decisions about a sensor, how long is the data valid. Finally we'd like to be able to display the live data on this schematic as the thing runs. We'd like to user to be able to enter the rules in a natural language, that is his language, if possible we'd like to be able to set up his system so the user doesn't have to type anything with numbers, and then maybe messages that he sends off to an operator or to an engineer or to a log. What we would just like the process engineer to do is sit down, use some kind of menu selection device and build up his rules from the grammar that is inherent into the rule entry system. These rules here can be very different kinds of rules and statements. All these came out, all these values here, reflecting conditions of tanks and valves, came out from a menu selection procedure so that the user doesn't have to type anything in. The only thing the user types in here is .4545, the number. The user doesn't type in anything else. Based on the information in the schematic down here, the system can generate the grammar that will enable the process engineer to go off and build his rules without the need to program, without being a knowledge engineer, with only knowing his job. And that is really handy for the process engineer because when its time for him to detect that a rule is wrong and to fix it up or add another rule he goes off and does it. He doesn't have to go off and call the programmer and hope the guy will get over there in a couple of days to help him out. So what we need are various kinds of rules, in the language of the process engineer. We have "IF THEN" rules, we have unconditional rules, for example if you are a discrete operation, such as a drilling operation, you may have step variables that define step 1, step 2, step 3, step 4 of the operations. You want the user to be able to enter those operations and set them up and what have you in his own language. Some people like to use "WHENEVER", some people like to set up initial values using initial statements and there are just a myriad of different kinds of rules that you can have and different kinds of statements you can have in order to meet user specification. Now the most exciting one here is the generic rule. You've got a system that has 40,000 pumps in it, all of the same kind, all of the same brand, all of the same pressure, you would like to be able to say, if anything happens to one of these pumps, I want you to tell me about it. What you don't want to say if anything happens to pump 1, I want you to tell me about it, if anything happens to pump 2, I want you to tell me about it and go through 40,000 of these things. You'd like to have a generic rule so that you can reference pumps in general. If the error code of any modem on any processor is on, then send the modem of this kind of processor a message. It turns out that in this application for the network analysis, we had hundreds of modems and hundreds of processors and we write one rule to cover them all. In doing that the user doesn't need to worry about all the separate modems out there, he states his problem in general. Now the nice thing about these rules, is that the user defines the rules in his own natural language as he understands it for his own industry, and he also selects the interface mode. Interface mode means how often you are going to look at the knowledge, when should you invoke the rule, what kind of an inference method you are going to use, backward or forward chaining or another mechanism. You remember throughout the day you probably heard the word frames, used a little bit, someone talked about frames this morning, which is a way to associate particular attributes with particular values for particular objects. What an ideal system has is in today's technology, is a frame based system in which rules and objects and everything is based on the frame, that results in a very efficient implementation. So if you look at the frame for IF THEN rules, what you would like to be able to do is to define a category "safety". Someplace in your program you may say, if part such and such is too hot, the temperature on a certain sensor is too high, then I want you to activate all rules with category safety and the system goes off and does that for you. That is in process engineer's language. You can tell the system how often you want to work at this rule, how often you want to fire it off and once you fire it off, if it is true if you fire off this rule and it turns out that the condition of this rule is true then after that I want you to check this rule every two seconds, because once the thing becomes true it may be critical that you check them more often. In a real time system you are able to do that kind of thing. In addition, you want to know who put this rule in here, when it was changed and what the status of the thing is. Now if you look at how you how rules are used, we call this the inference paradigms. We have the traditional ones that people talk about all the time, in expert systems there is forward chaining, there is backward chaining, there is a couple of other ones that you need in a real time system. You need to be able to focus on particular portions of the system that are having their problems. You don't want to focus, you don't want to fire off all these rules all the time, you only want to focus on those particular rules when a particular condition occurs. The other kind of inference paradigm you want to use is that sometimes you want to go up and check rules just because its time to check, thats called the scan, especially in an industrial world a program has a scan in it, start at the top and go to the bottom, after a particular scan time, how long it took to go through there. The PICON system gives us a lot of ways to go off and check rules without the data having been flagged. In a system that is strictly backward chaining, or forward chaining, you need certain conditions to occur before you can fire those rules off. Suppose you want to go up and check those conditions before they occur. You want to do some "WHAT IF" check on it. Finally, if you look at the inference in this system, there is forward chaining, backward chaining, there is focussing, you can have this thing, and this is critical, the operator has to be able to have explanations on the conclusions from the various chaining mechanisms. You have to be able to go back and say why did you want to know that, why did you send that message, what is really going on here, and in a real time system such as PICON, this is a key thing. This is an example of building up a rule. You start off in building a rule by clicking on a menu and this starts it off for you. It says IF and then you take your mouse and you put it down on a sensor, that sensor up over here, and you continue selecting the elements of your rule by clicking on menu items. Do you understand the difference between forward chaining and backward chaining? One word: forward chaining means I have a set of data I need to make a conclusion. I have a bunch of data here I have to make a conclusion. Backward chaining is I have a problem, why do I have that problem. Thats the two basic differences between forward and backward, forward starts with data and ends up with a state, backward starts with a state and figures out why and does a diagnosis on it. The focus method is little different than any of the others and in order to do that you have to tell the system what to focus on. You can put a rule together and you can say that when a particular condition occurs, I want you to focus on furnace no. 3, and safety so that anything associated with this furnace and has a safety category will be fired off. Now we've got this database out there, we have done the knowledge acquisition, built up our rules, built up our schematic, we have worried about the things we needed to worry about in real time and we've got 100,000 rules out there, we've got 200,000 objects, and we want to go off and find out all the objects that reference FR cell; all the rules that reference FR cell. We want to find out all the sensors that have a scanner level greater than 30. How do we do that? Well in a large expert system environment, you need to have a place to examine your knowledge base and in the one we are talking about here there is a relational retrieval mechanism, where you can retrieve just about any information that you have in your database, in your knowledge base, with simple rules like this. You can say retrieve all rules associated with safety, retrieve all rules mentioning component XYZ, retrieve rules having problems. What we can say here is, what we can do with this is, we don't want, first of all, we don't want the process engineer to have to type anything. We don't want the process engineer to make any mistakes typing, we don't want the process engineer, any engineer or the user or whoever it is to ask for things that aren't there. So what we do is, we come over and I click on something that says retrieve and I say retrieve parameters for knowledge base and rules such and such. The rest of this stuff I did simply by clocking on various things, in this random selection machine in here. When I'm all done I go over here and hit the end key, at that point it goes off and retrieves everything in the system. The last thing we want to do is look at all the kncwledge in there and get it checked out before we put it on line we want to simulate it, so what we need it a way to mathematically describe our system to the expert system, we need to mathematically describe our process to the expert system, and we can do that with something called a simulation statement. Once we get our simulation under control, we can go off and do on line testing and the way you do on line testing is somehow tapping data out of your system and feed it into your expert system, the real time system, so that that expert system thinks its actually on line. And then the last thing you do after your data testing is all successful is to put it on line. You never put this thing on line until you've been successful with it. This is the simulation running, it tells you the values of all the sensors you have on here, and also reports all the messages that have come back from the operator.

Now, I'm not going to go through the rest of the slides I have, but to summarize it, I believe that the future of expert systems is real time, that is what we've got to get, we have to work out all the problems involved in expert systems, in real time expert systems and in real time which is separate from expert system. And more than that, we have to worry about how we are going to get the end user develop these systems, because we can't afford to have three people working on every expert system we have, one expert, one programmer and one knowledge engineer, sitting around trying to get this thing to go. What we need are expert system environments that run in the real time world to enable the user to sit down and bring his knowledge into the system and run it and test it and see if he comes out with the same answer. And then finally, one point that nobody is addressing right now: how you know that data base is right? I started off talking about an expert system, that had 100,000 frames in it, several hundred thousand rules, several hundred thousand objects, do you mean to tell me there are no bugs in there? This thing is controlling an entire plant. To bring it closer to home, that expert system is controlling the automatic money machine that you stick a quarter in and you just punch the buttons, say give me some money, transfer the money out of your bank account. It worked, the expert system said it is ok. How do you know these things are not coded false. This is a totally different subject, but when you analyse the system you got to worry about that. Thank you.

PC BASED A.I. SYSTEM DEVELOPMENT TOOLS

T. GOMI,
APPLIED AI SYSTEMS INC.,
KANATA

NOTE FROM THE EDITOR:

It was not possible to transcribe the recording of this presentation. The author has provided a copy of all the visual aids used, several of which are self-explanatory. They are reproduced as received.

SMALL TOOLS

- What's happening in PC-based AI Tools
- New Generation PC-based AI Tools
- Why PC-based Tools?
- Some Example PC-based Tools
- Future!

SMALL TOOLS

- Analysis/Observations
  - Users realized the ease of use of a simple rule-based, goal-driven, back-tracking only shells of 1985
  - At the same time, they grew tired of the limitations of such tools in:
    - Developing practical size KBs
    - Expressing various knowledge types
    - Properly structuring KBs
    - Making flexible user interface
    - Accessing various sources of real data
  - Some ESs actually discredited AI

Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Analysis/Observations (Cont.)

  - Limitations in user interface the most critical issue in many practical application:

    - And no tools were solving them

    - Nobody is going to answer 20 questions, only to get an "Yes"

  - Inability to cooperate with other computational models:

    - 2GL: assemblers for special input/output

    - 3GL: Fortran, C, Pascal, Ada, etc. for necessary procedural knowledge representation

    - 4GL: Databases, spreadsheets, querry processing languages

SMALL TOOLS

- Second Generation Small Tools
  - No longer a pure declarative programming tool
  - Hybrid computational models in a number of ways
  - Hybrid Knowledge Representation:
    - Rules or predicate calculus
    - Frames, schema, or objects
    - Procedural knowledge:
  - Hybrid reasoning:
    - Backward chaining
    - Forward chaining
    - Procedural attachments
    - Taxonomical reasoning

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Second Generation Small Tools (Cont.)

  - Some shooting for more capable processors

    - 80286, 80386
    - Hypercube

  - Language interface:

    Many now can talk to 2nd, 3rd, and 4th generation languages

    Some early tools had the linkage capability, but extremely hard to use. Most of such problems are solved:

- The difference between large tools and the small tools becoming fuzzy

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Second Generation Small Tools (Cont.)
  - Language interface (Cont.):
    - "Any language produced for MS-DOS by Microsoft"
    - Common runtime support among languages the key for versatility
    - Some still require considerable engineering level effort to link up
- IBM PC and Macintosh enjoy the vast majority of these tools

  Tools very slow in development or non-existent for other small computers

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Second Generation Small Tools (Cont.)

  - Ability to add meta-knowledge. Typically, in the form of rules in a separate KB

  - Improved natural language frontending (both input and output)

    Still a template matching (input) or a canned phrases (output), but greatly more flexible, editable, and modifiable

  - Access to many amenities:

    - Color graphics
    - Multiple windows
    - Mouse and other pointing devices
    - Menues

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

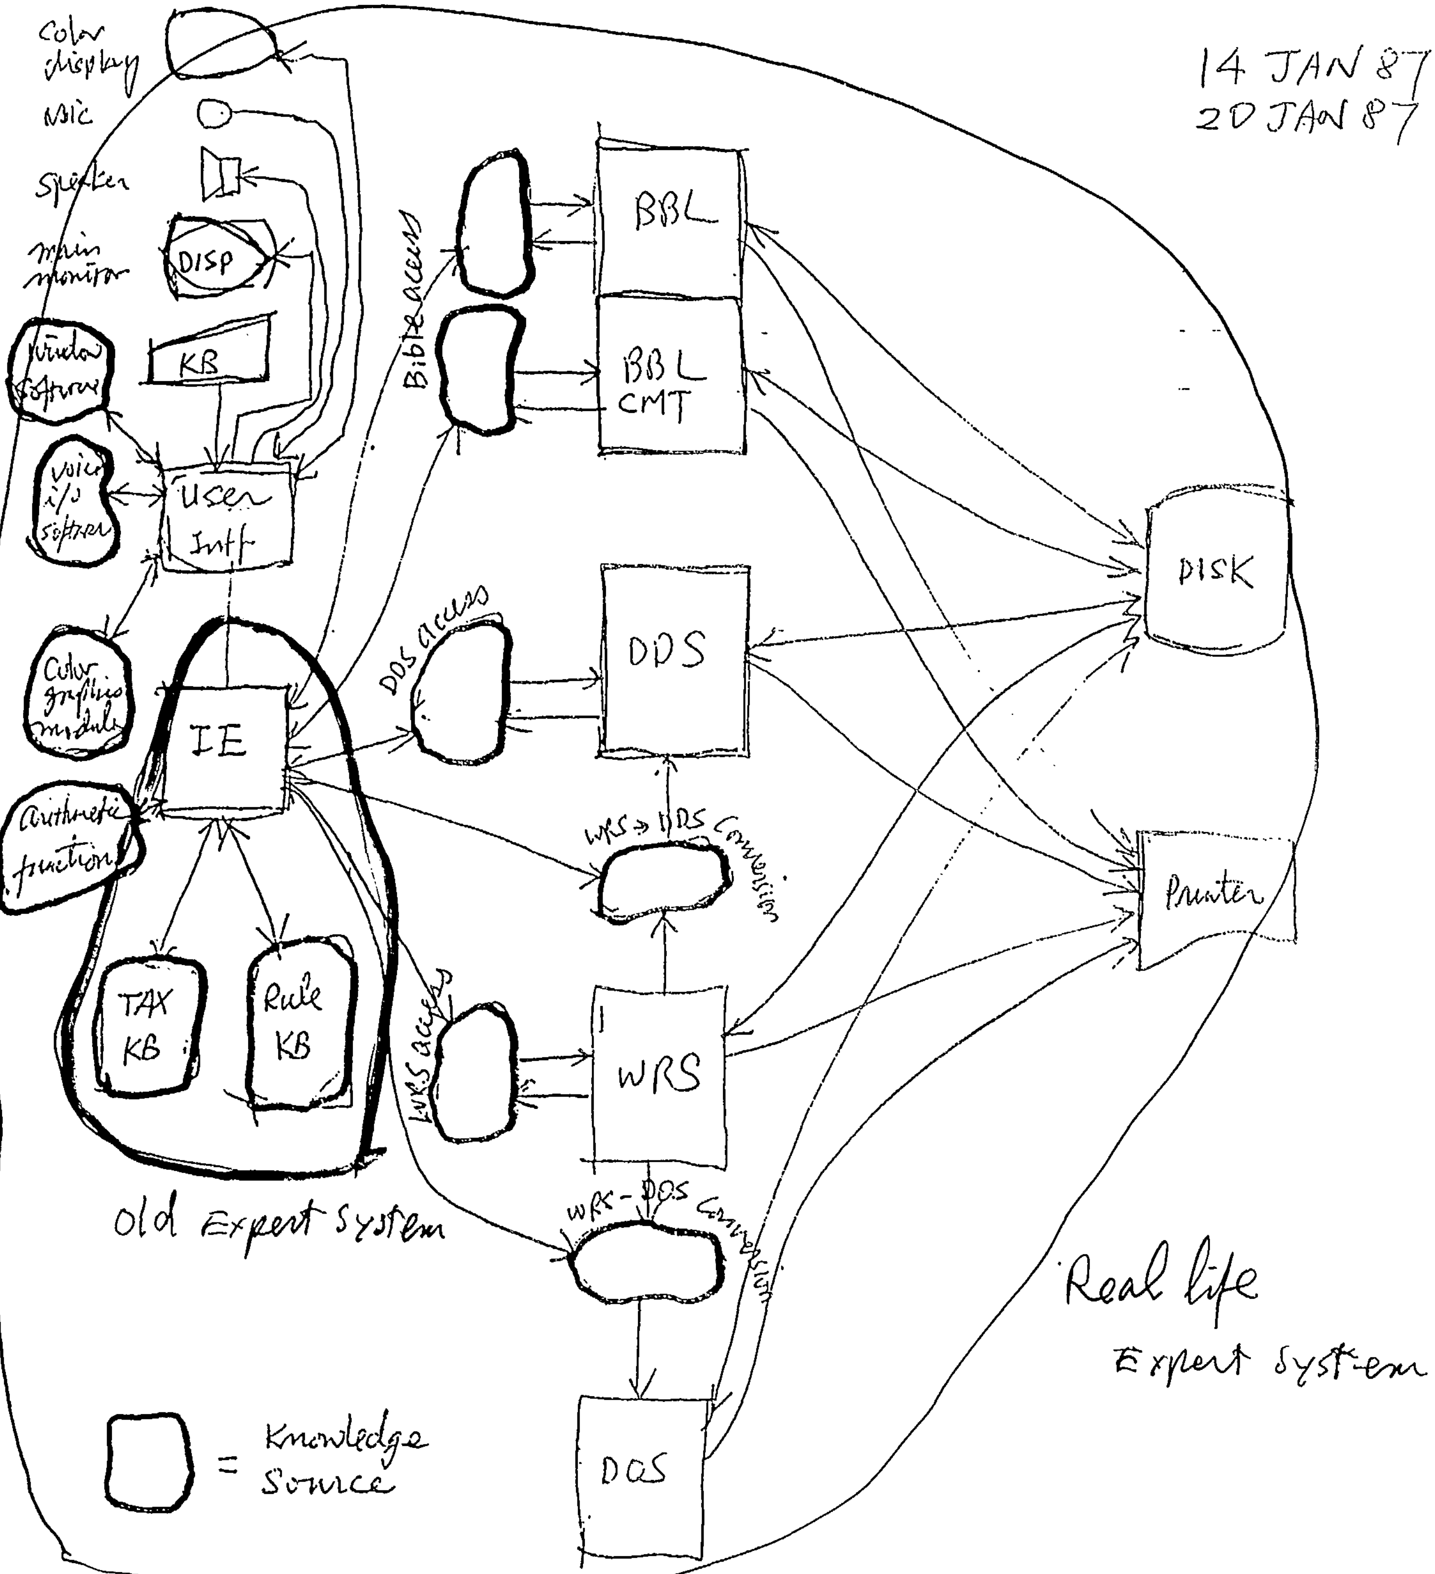

2

Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

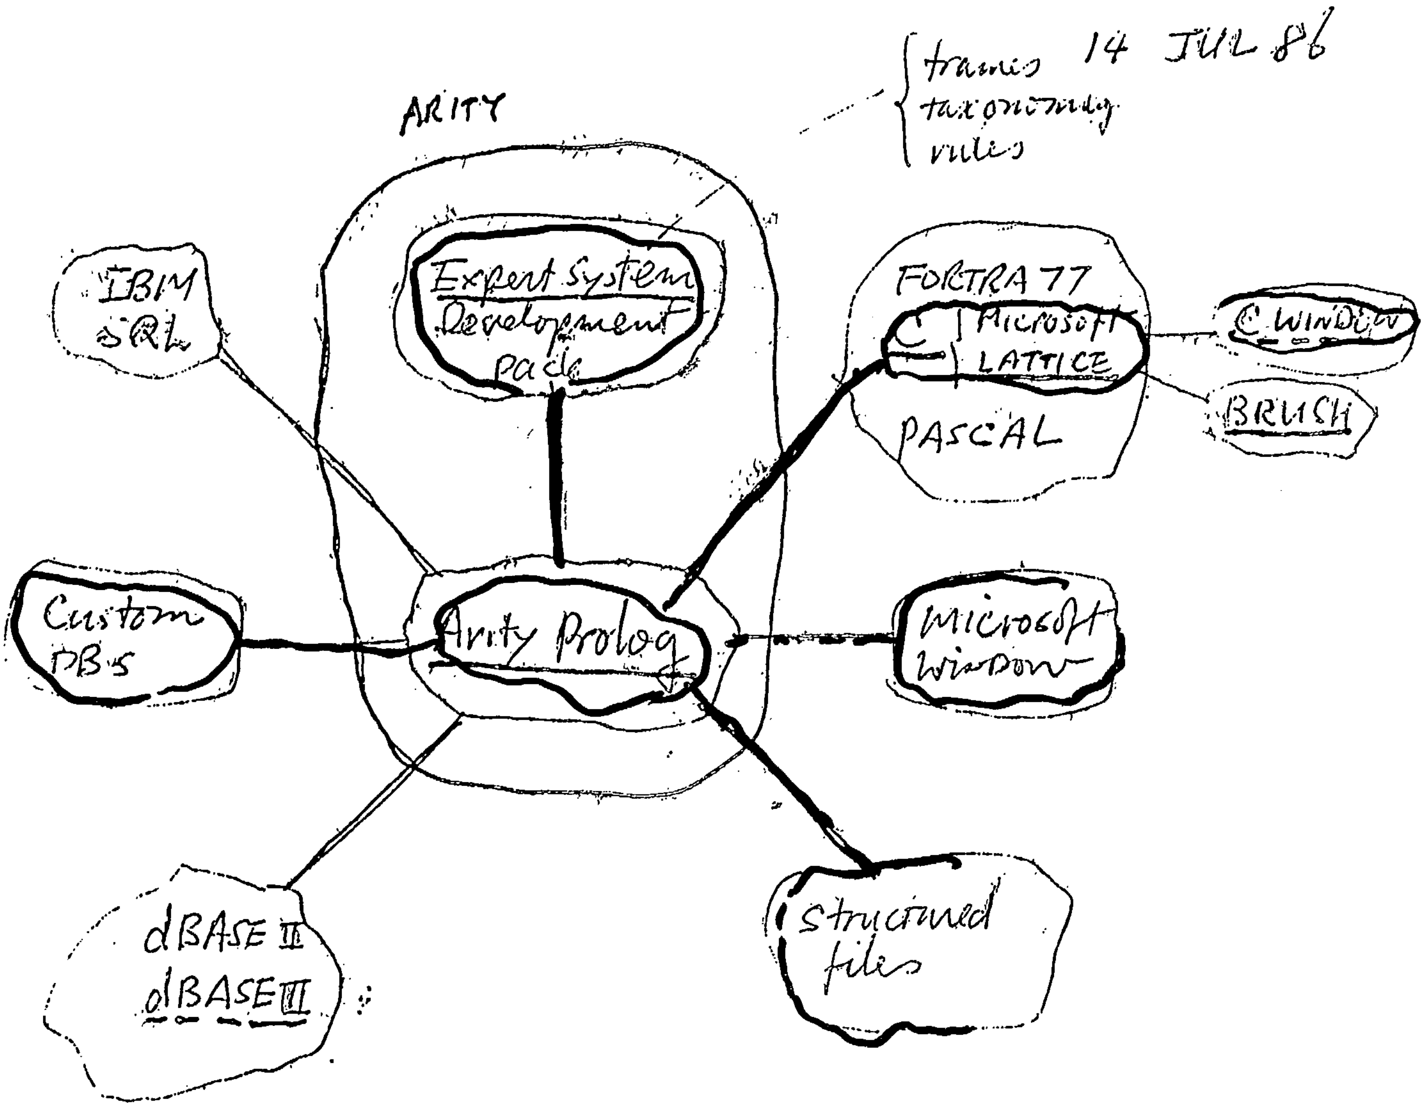

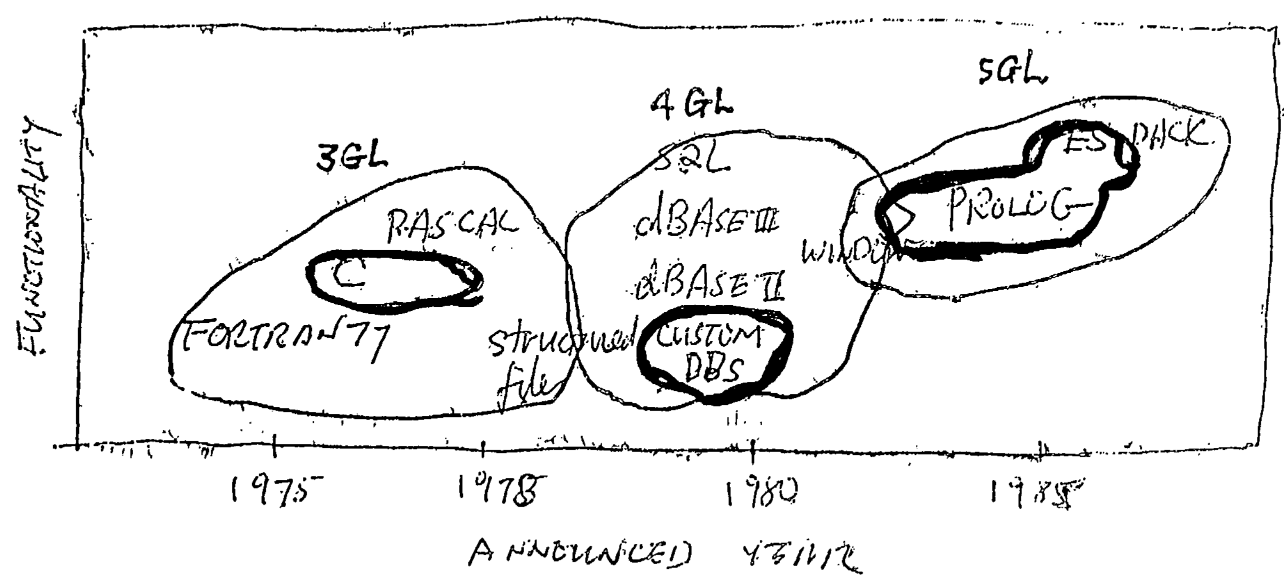

Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

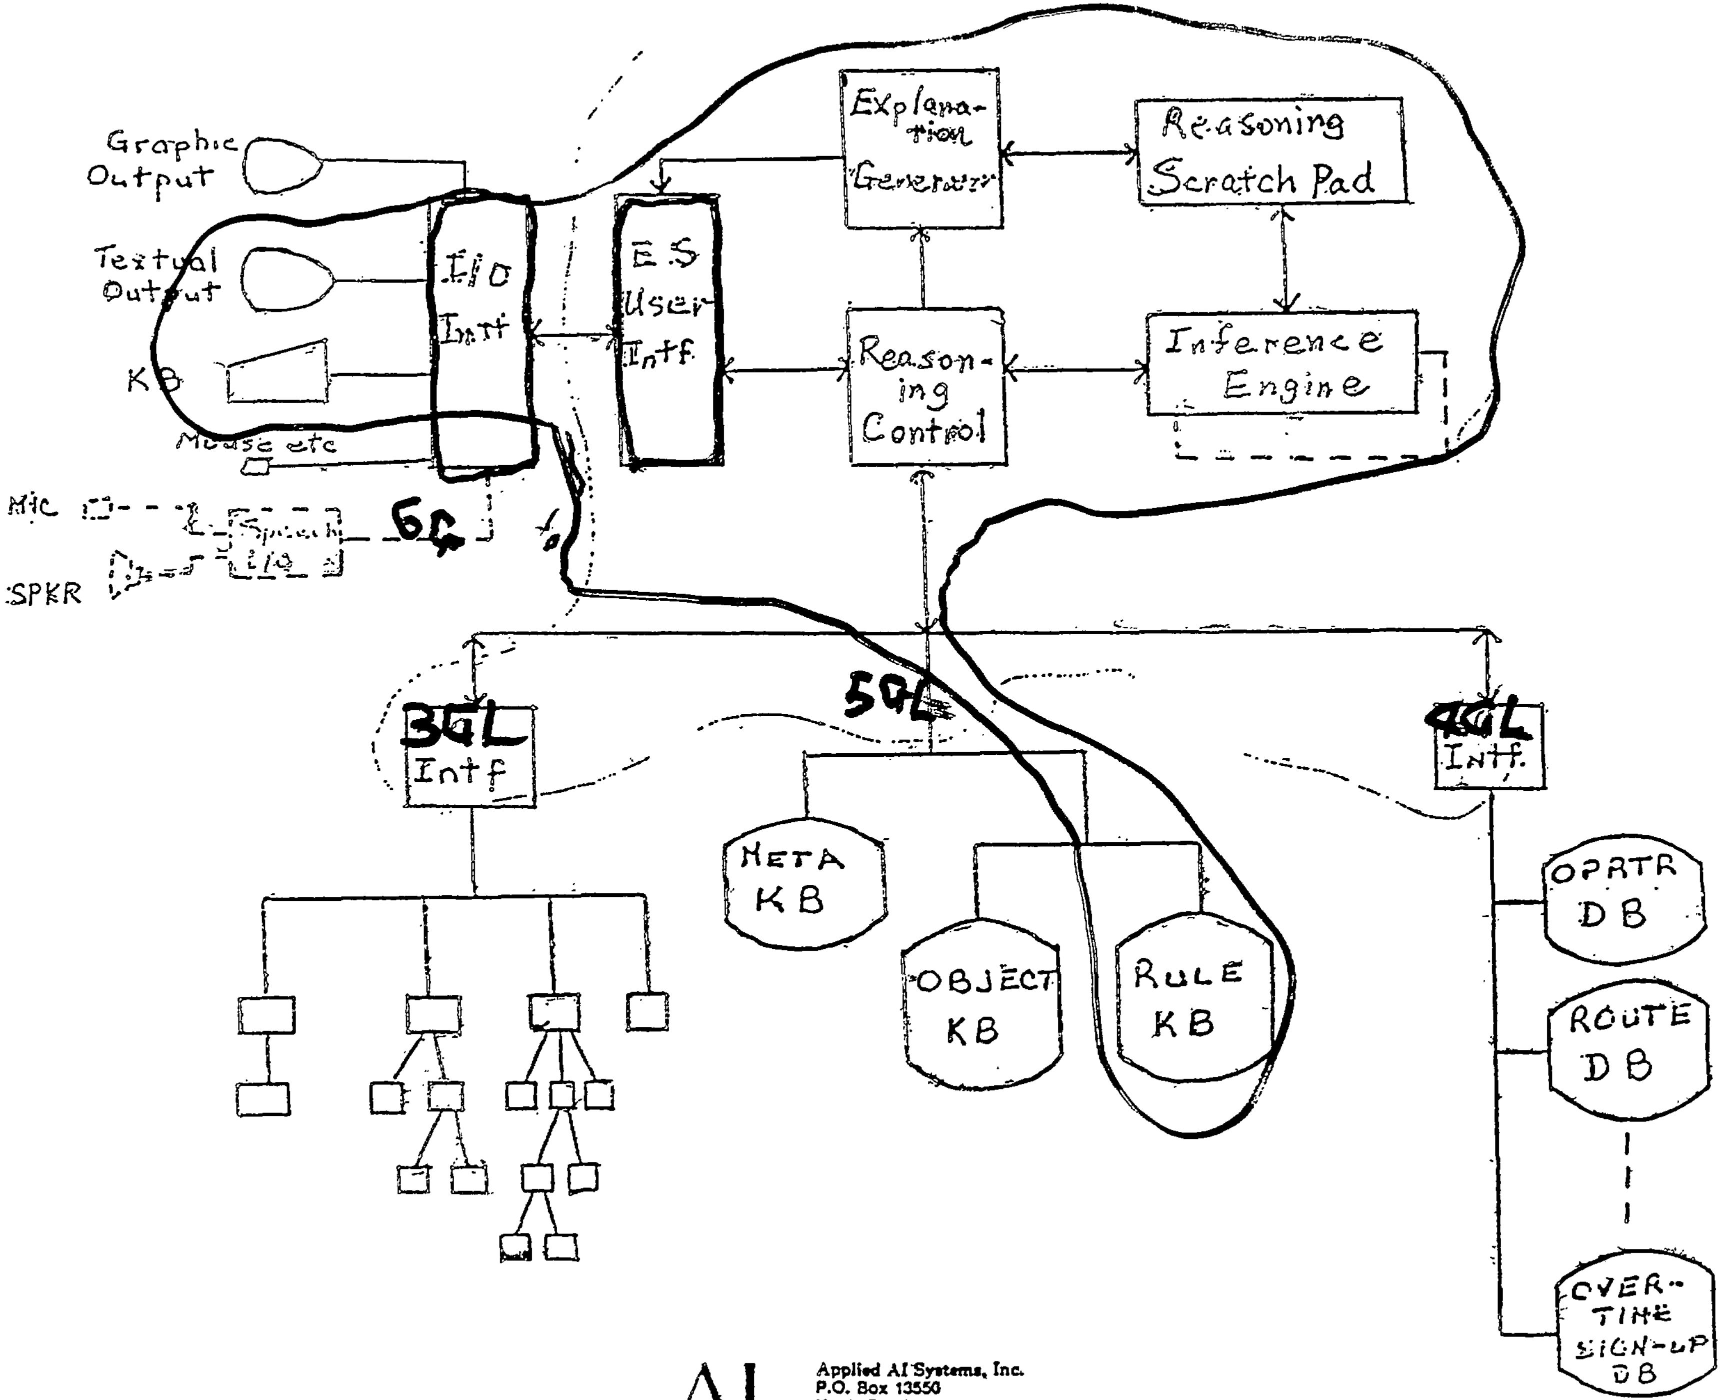

Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X8
Telephone: (613) 592-0084

SMALL TOOLS

- Examples
  - Xi
    - IBM PC class machines
    - Rule plus frame
    - Extensive explanation facility
    - Menu interface
    - Adjstable natural language frontend
  - GURU
    - IBM PC class machines
    - For "business applications"
    - Both backward and forward chaining
    - Automatic solicitation of input
    - Excellent DB access
    - Handling of uncertainty
    - Access to spreadsheet
    - Business graphics
    - 3GL linkage
    - Menu interface

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Examples (Cont.)

  - micro-PROLOG/APES

    - For IBM PCs and Macintosh
    - For scientific applications
    - One of the most elegent implementations
    - Limited linkage to 3GL
    - Template NL interface
    - Automatic solicitation of input
    - Canned phrses for explanation
    - Windows
    - Menues

  - Insite

    - IBM PC class machines
    - Still rule-based
    - Access to DB
    - Linkage to 3GL
    - Windows
    - Menues

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Examples (Cont.)
  - Arity
    - IBM PC class machines
    - For "business applications"
    - Excellent DB access
    - Access to spreadsheet
    - Handling of uncertainty
    - Business graphics
    - 3GL linkage
    - Menu interface
  - MacScheme
    - For Macintosh
    - Object-oriented Lisp extention
    - Ability to define 3GL constructs
    - Extensive macro capabilities

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Examples (Cont.)

  - GC Lisp

    - For IBM PCs
    - Good implementation of Common LISP standard
    - Powerful ES shell coming
    - Graphic package
    - Built-in LISP tutor
    - Version for 80286
    - Aiming at a powerful AI workstation using 80386

  - Smalltalk/V

    - For IBM PCs
    - Inexpensive but powerful
    - Multiple windows
    - Object-oriented
    - Rules and class hierarchy
    - Bitmap graphics with editor

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

SMALL TOOLS

- Examples (Cont.)

  - Q&A

    - For IBM PCs
    - Intelligent DB with natural language front end
    - Relational DB
    - Can process standard DB data (e.g., dBase III)
    - Can process spreadsheet data

Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

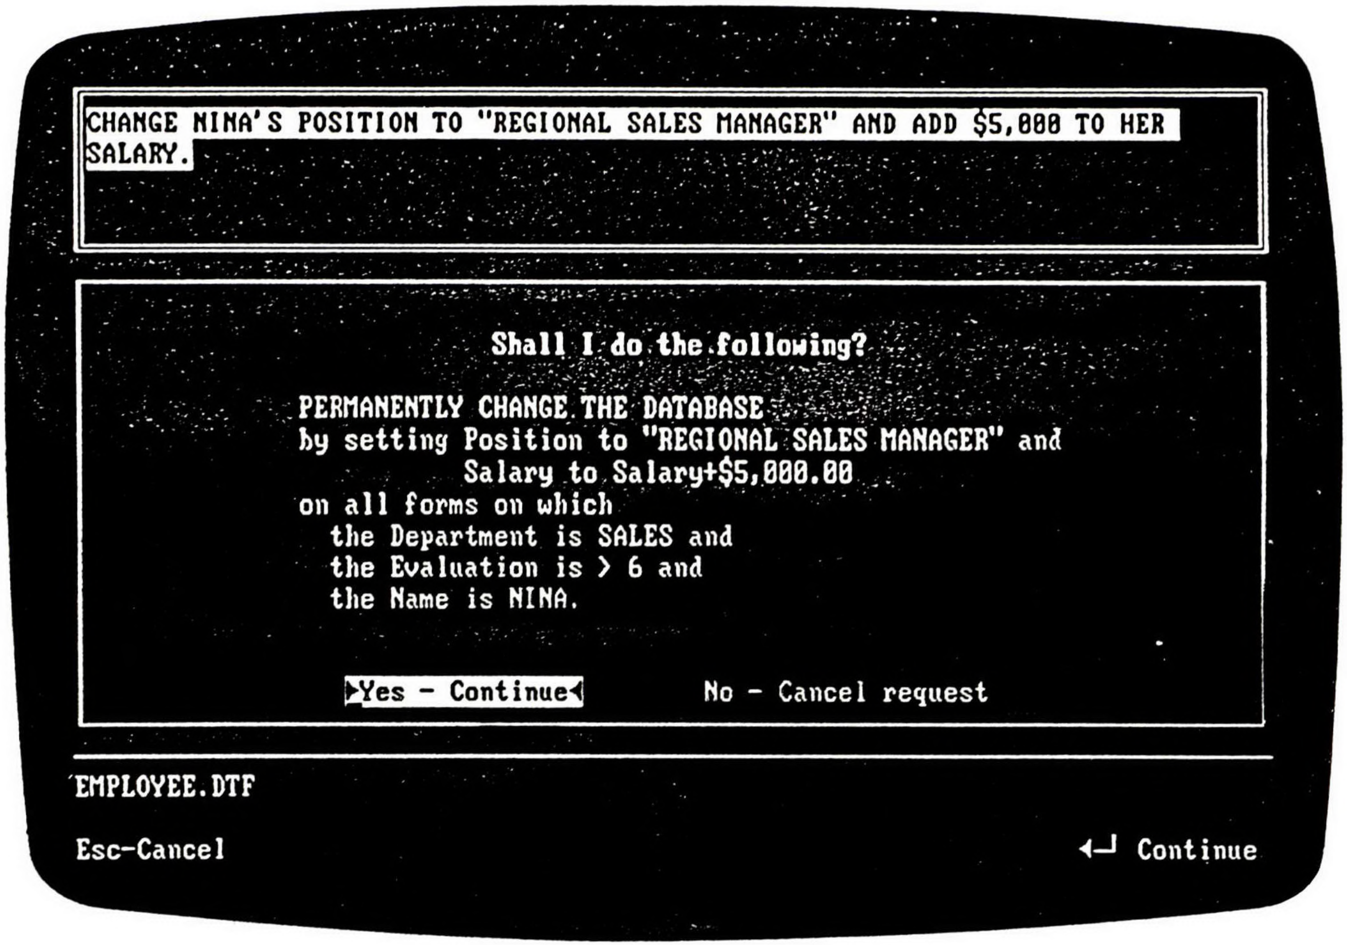
CHANGE NINA'S POSITION TO "REGIONAL SALES MANAGER" AND ADD $5,000 TO HER SALARY.
Shall I do the following?
PERMANENTLY CHANGE THE DATABASE
by setting Position to "REGIONAL SALES MANAGER" and
Salary to Salary+$5,000.00
on all forms on which
the Department is SALES and
the Evaluation is > 6 and
the Name is NINA.
Yes - Continue
No - Cancel request
EMPLOYEE.DTF
Esc-Cancel
Continue

SMALL TOOLS

- Why PC-based Tools?
  - Unquestionablly more cost effective
    - Lisp machines failure Cause #1
  - Over 95% of time one is developing software, not running it
    - Human-interactive speed, not Lisp machine's crunching power
  - Smoother transition to System Delivery
    - What you build is what you run
    - Lisp machine failure Cause #2
  - One can begin AI today

- Why PC-based Tools? (Cont.)

  - Cooperation with over 50,000 off the shelf software

    - Hybrid system key to practical AI applications

  - No need to "sell" hardware
    - PCs everywhere
    - Some awaiting more usage
    - Classes of PC tools available

  - Powerful PCs rushing in
    - 386 PCs
    - Macintosh II
    - Hypercube parallel machine
    - IBM's announcement later this year

6. E.A.S.T. EUREKA ADVANCED SOFTWARE TECHNOLOGY

1. CONTENT

Development of manufacturing cells workshop on software engineering related to UNIX system and EMERAUDE system, the industrialized version of ESPRIT PCTE prototype. The product covers many activities such as computer system, business applications and artificial intelligence.

2. PROJECT TIME

Six years

3. PROJECT COST

$150 M

4. PARTICIPATING COMPANIES/COUNTRIES

Switzerland (CIR)
Denmark (CRI Computer Resources INT)
Finland (Nokia)
Italy (Datauat, Intecs, Sesa, Italia, Selenia)
France (SFGL)
Great Britain (CAP Industries Ltd, 143-149 Farringdon Road, London EC1R3AD)
Spain (Sereland)

12. EUROPEAN SOFTWARE FACTORY

1. CONTENT

Design and creation of a database accessible to firms engaged in software development

2. PROJECT TIME

8 Years

3. PROJECT COST

$467 M

4. PARTICIPATING COMPANIES/COUNTRIES

France (Cap Gemini Soget)
FRG (Nixdorf) Krupp-Atlas, Elektronik GmbH, AEG, Softlab, IABG.
Norway
Spain (Sereland)
Sweden

Interest: Ireland, Commission

Fundamentals of Telephone Interpretation

—Pulling Together the Elements of Realtime Machine Translation over Telephone Lines—

Smooth communications between people of different countries speaking different languages—this has been one of mankind's dreams for many years. The day is soon coming, however, when there will be a real need for automatic translation and transmission of what is being spoken by parties connected via telephone lines. Basic research is accordingly going forward in regard to machine recognition of what a speaker says, translation of one language into another, computer speech synthesis, and related techniques. Model system tests are also being undertaken that will allow evaluation of systemized combinations of these techniques.

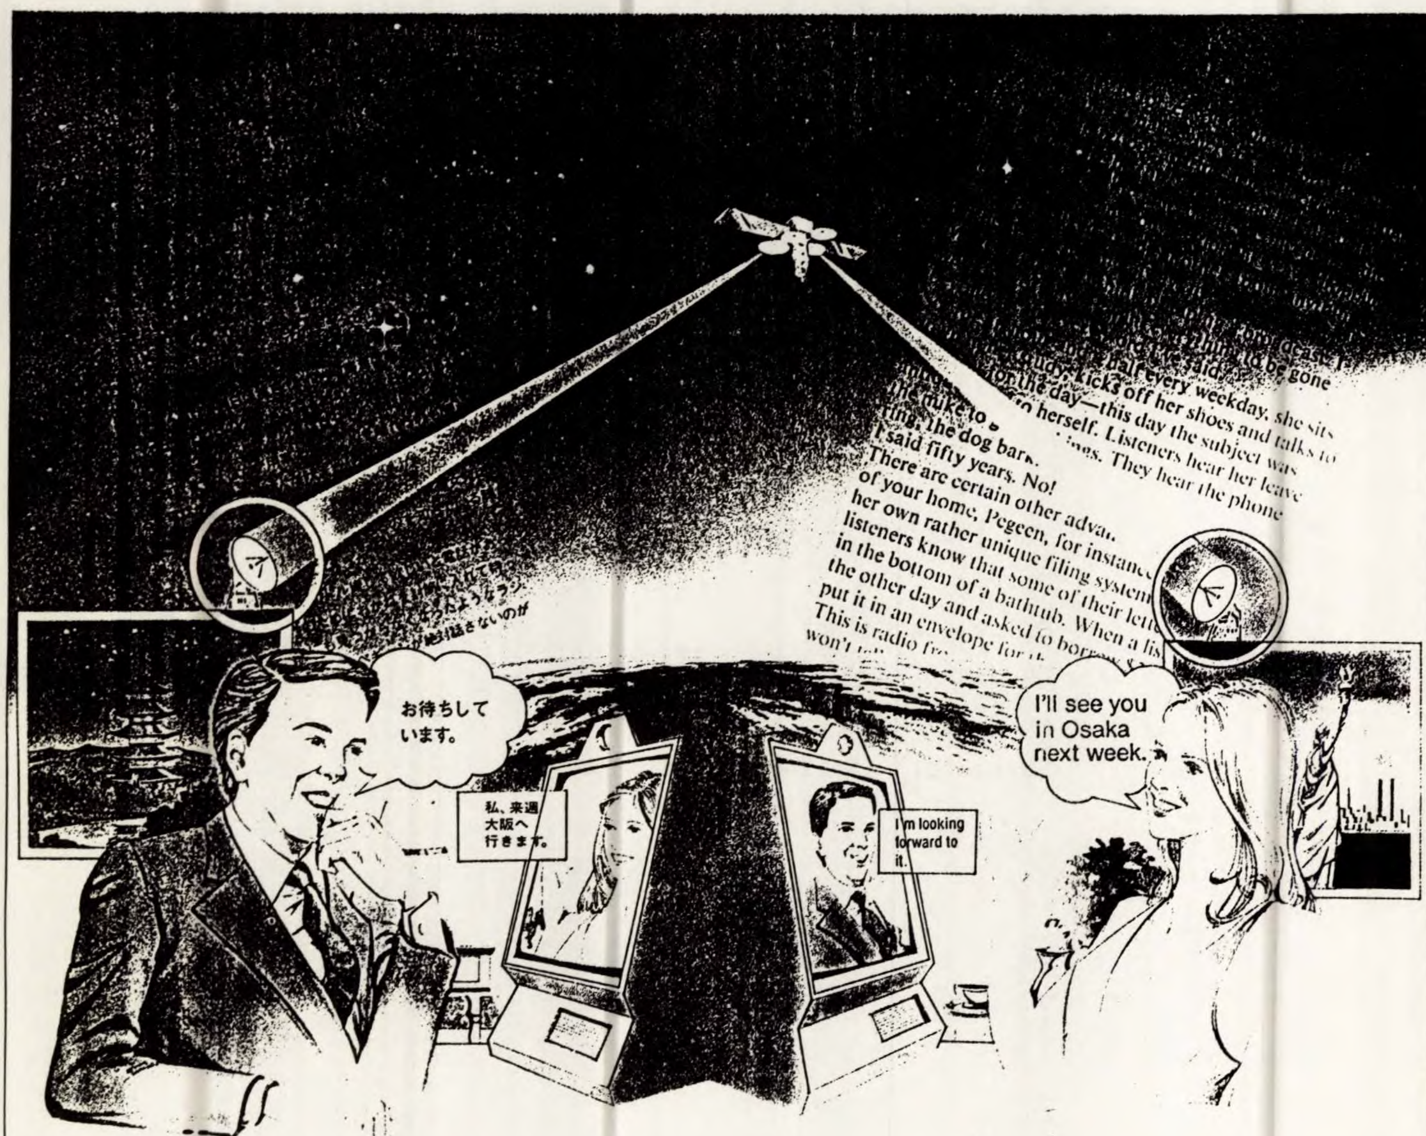

Human Science Fundamentals Involving Auditory and Visual Mechanisms

—Basic Research into Areas of Sight and Hearing Aimed at Better-engineered Human/Machine Interfaces—

Psychology, physiology and engineering are only some of the areas involved in the attempt to realize information processing and communications systems that are truly suited to human characteristics and that anyone can use easily. Work is proceeding in these areas toward modelling of perception and recognition processes in human sight and hearing, as well as of human thought process, learning and behavioral mechanisms. Aiming not only at major advances in character, figure, image and speech recognition technologies, progress is also targeted for clarification of the optimal form human/machine interfaces should take in the information-rich society we are currently evolving towards.

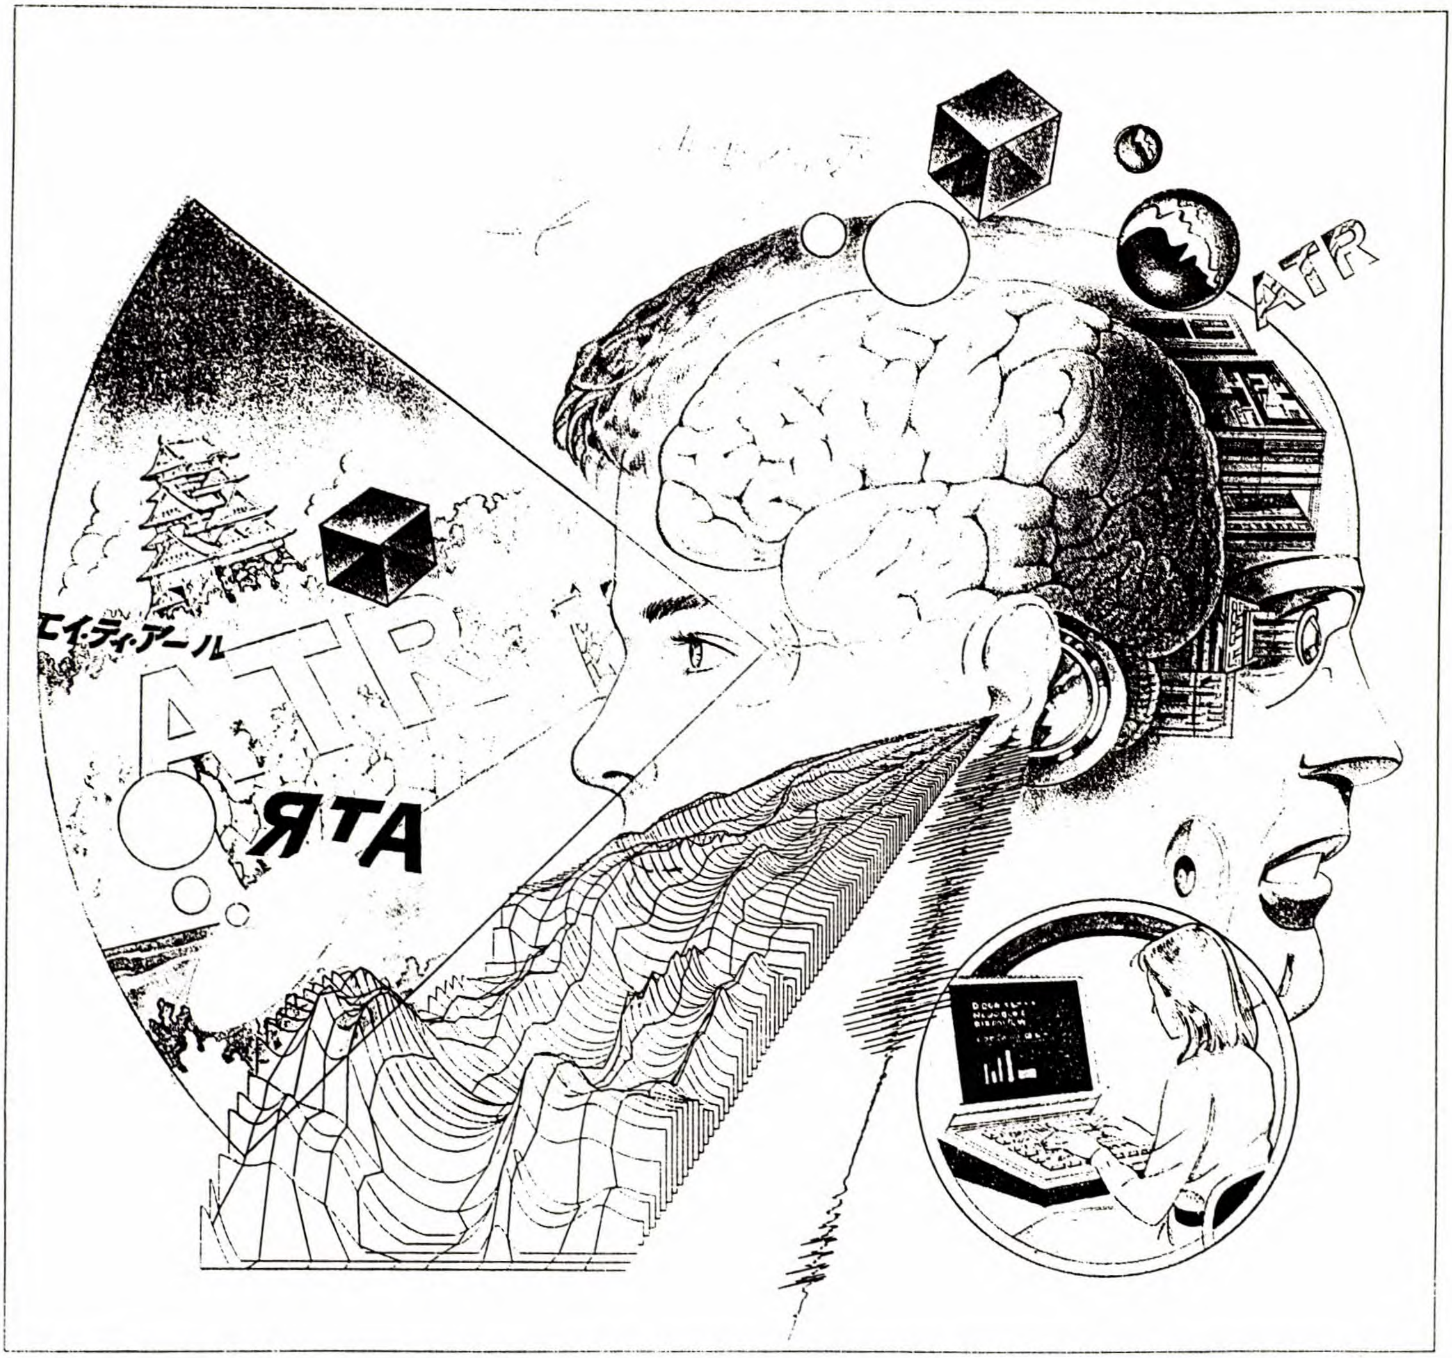

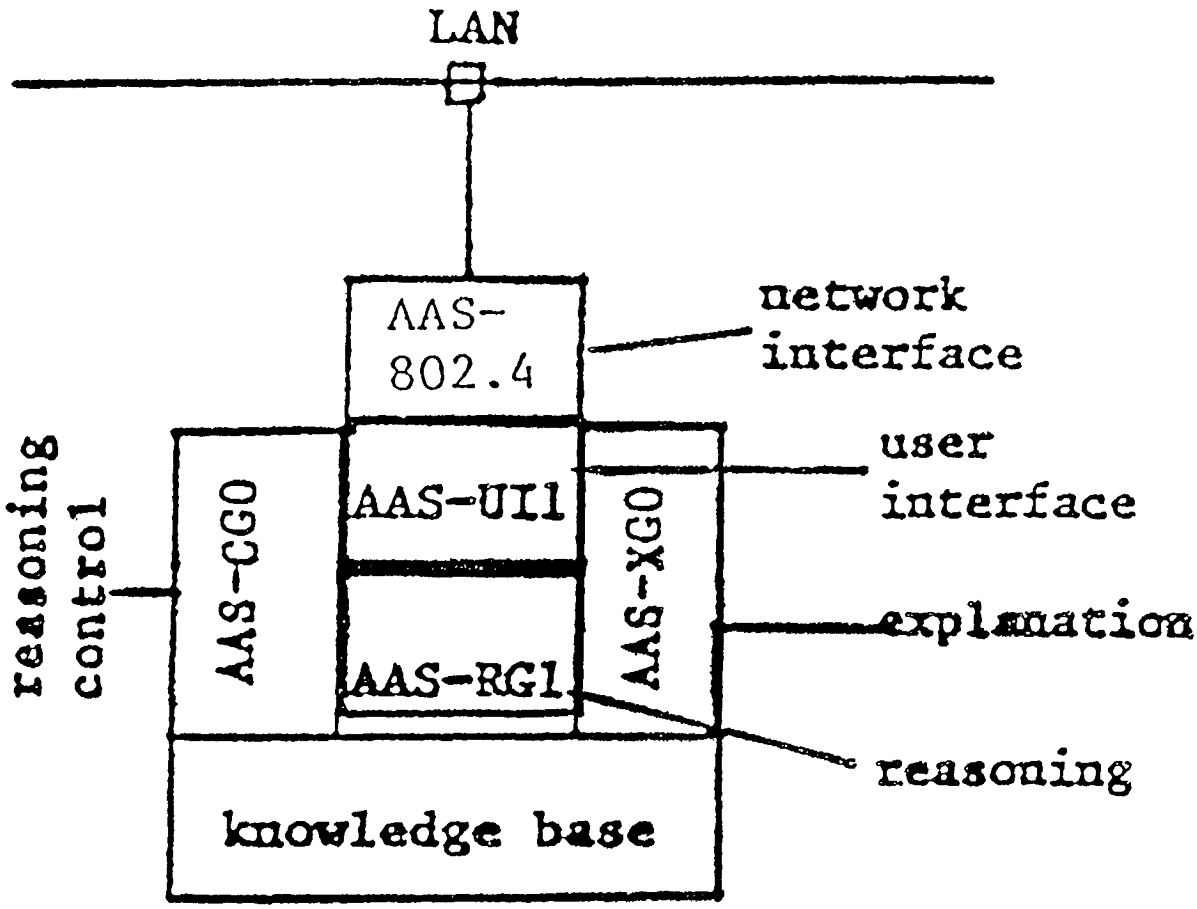
LAN
AAS-
802.4
network
interface
reasoning
control
AAS-CG0
AAS-UI1
user
interface
AAS-XG0
explanation
AAS-RG1
reasoning
knowledge base

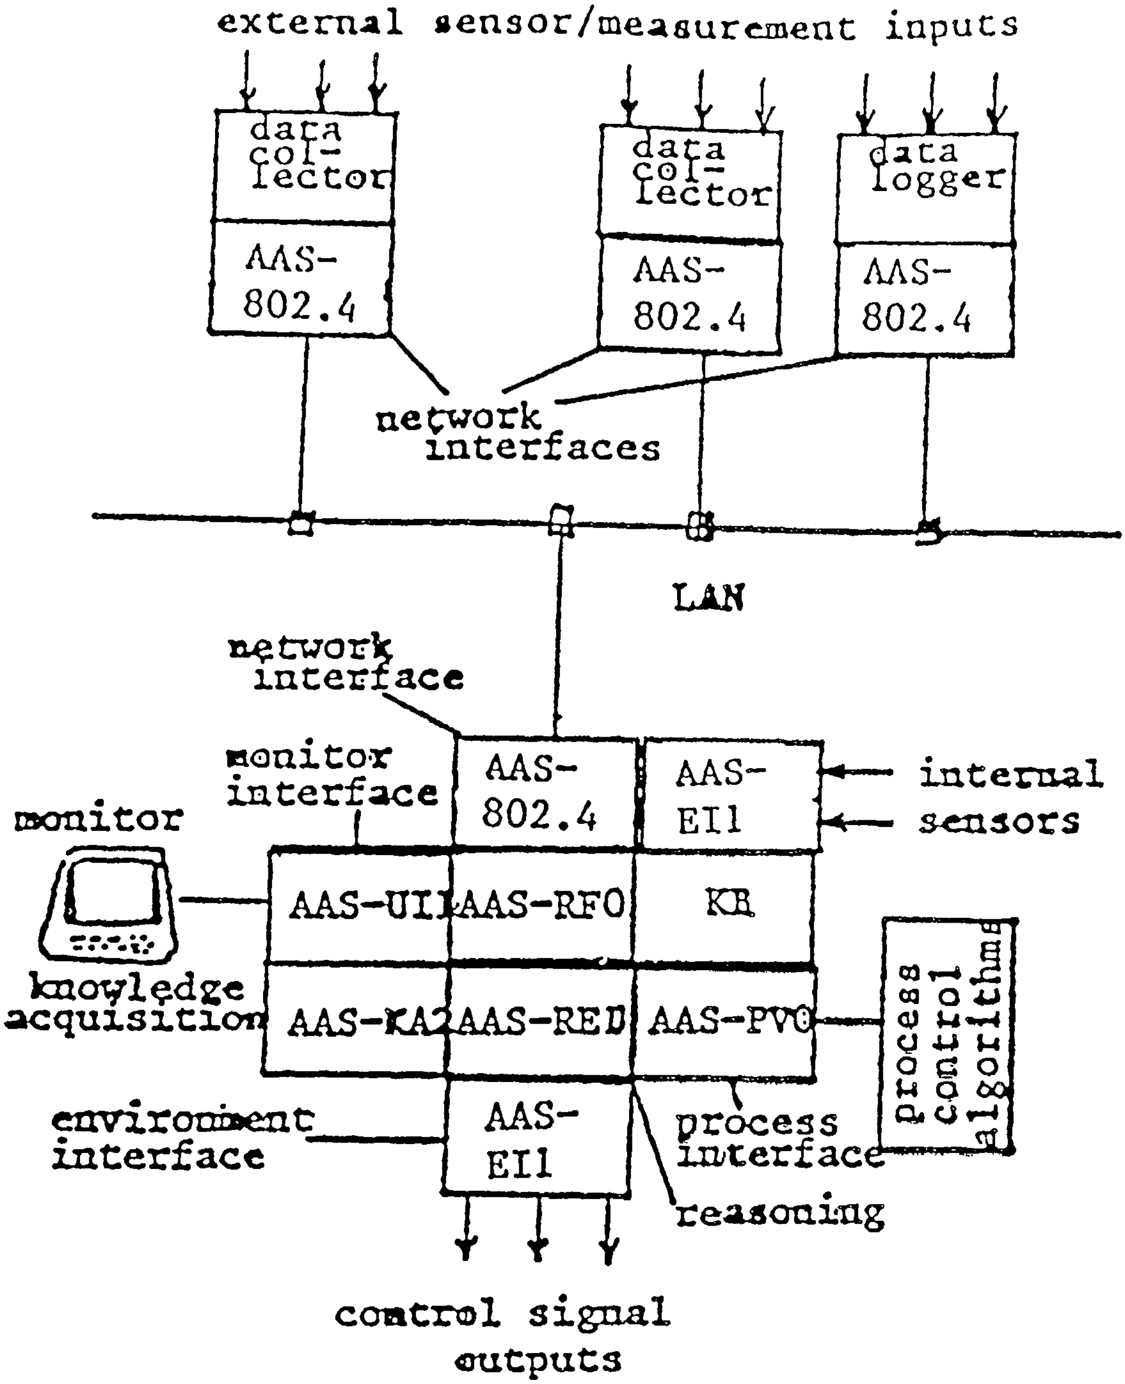
external sensor/measurement inputs
data collector
AAS-802.4
data collector
AAS-802.4
data logger
AAS-802.4
network interfaces
LAN
network interface
monitor interface
AAS-802.4
AAS-EI1
internal sensors
monitor
AAS-UI1
AAS-RFO
KB
process control algorithms
knowledge acquisition
AAS-KA2
AAS-RED
AAS-PVO
environment interface
AAS-EI1
process interface
reasoning
control signal outputs

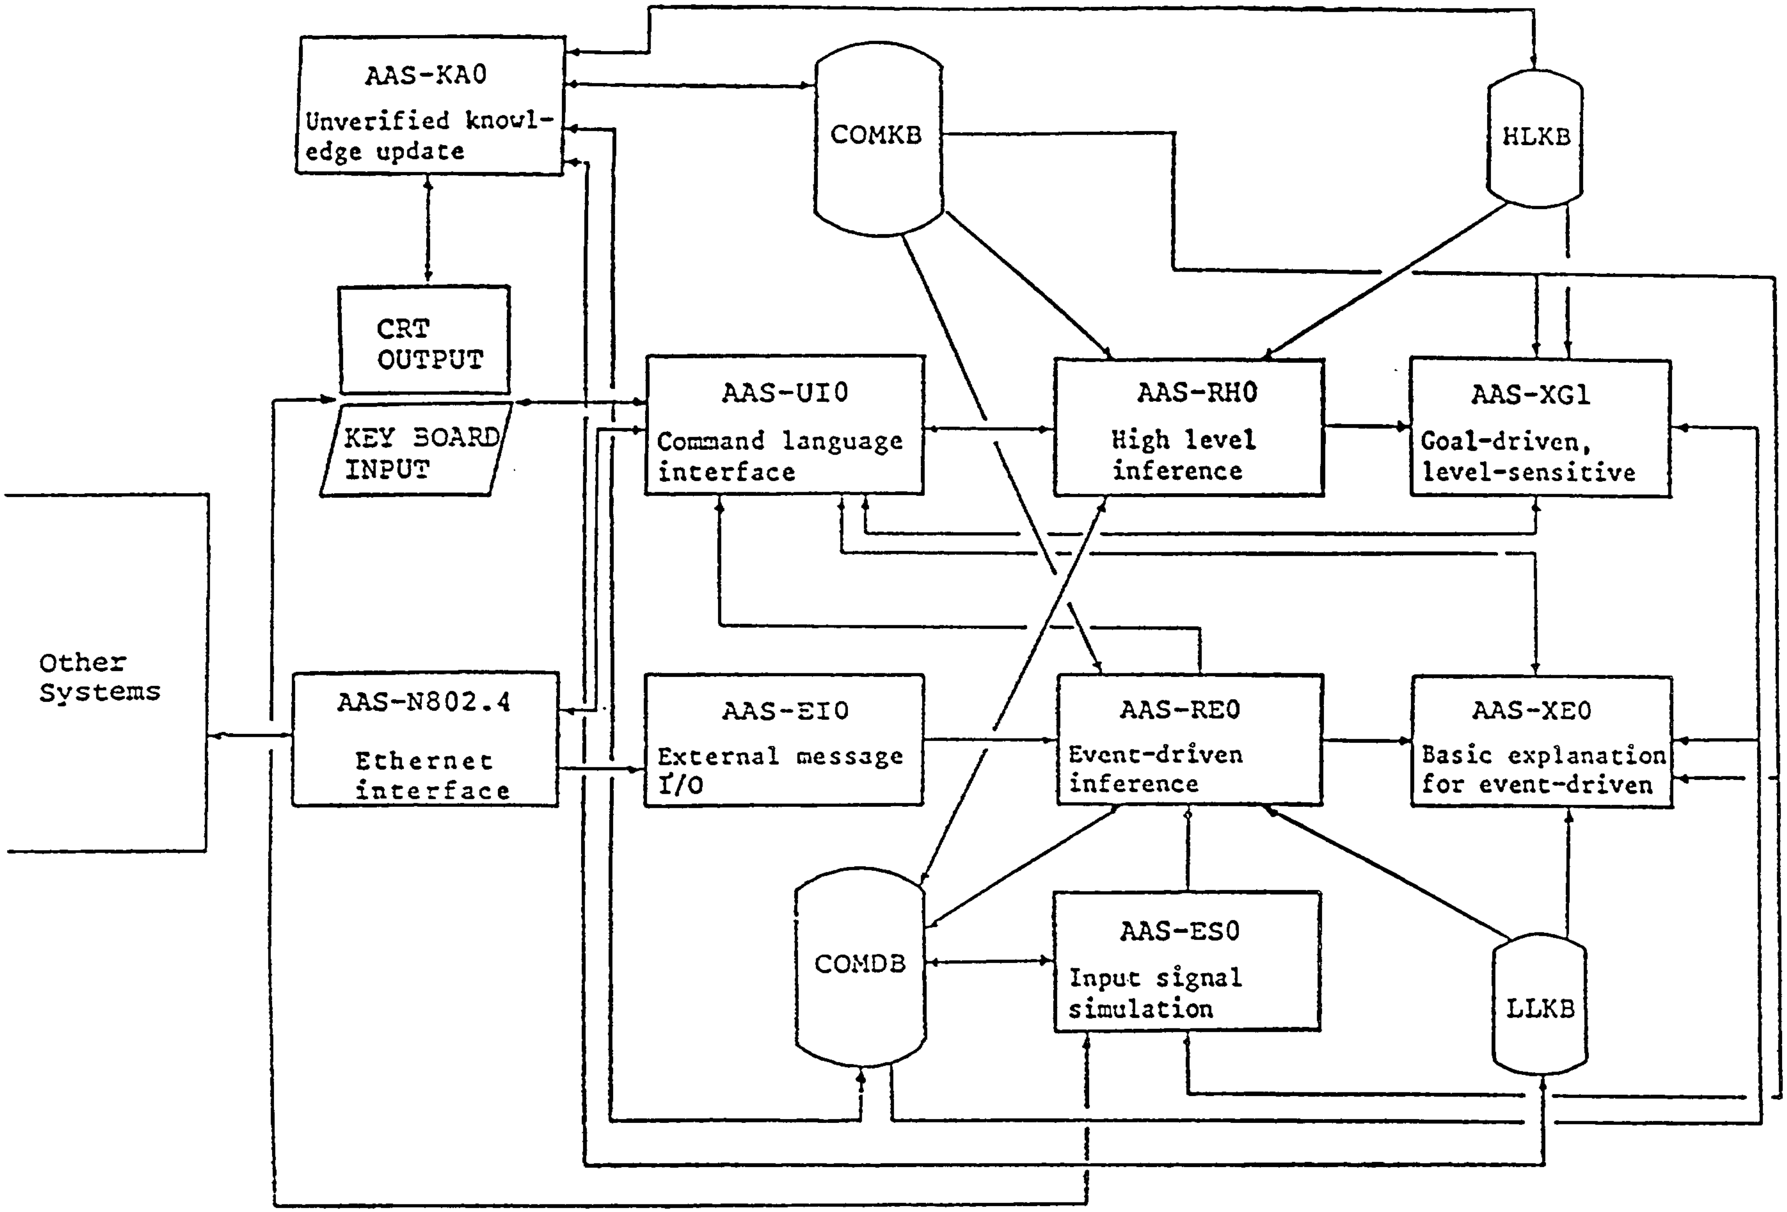
AAS-KA0
Unverified knowl-
edge update
COMKB
HLKB
CRT
OUTPUT
KEY BOARD
INPUT
AAS-UI0
Command language
interface
AAS-RH0
High level
inference
AAS-XG1
Goal-driven,
level-sensitive
Other
Systems
AAS-N802.4
Ethernet
interface
AAS-EI0
External message
I/O
AAS-RE0
Event-driven
inference
AAS-XE0
Basic explanation
for event-driven
COMDB
AAS-ES0
Input signal
simulation
LLKB

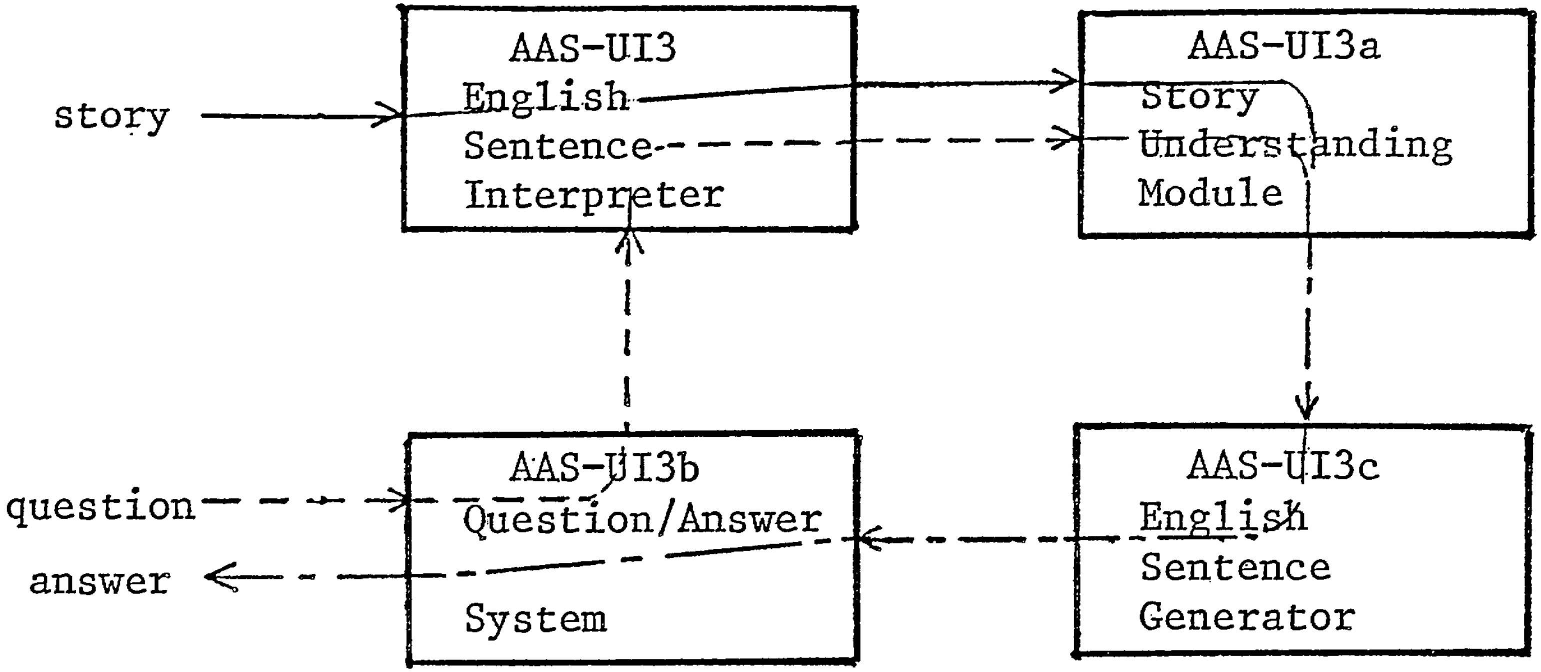
story
AAS-UI3
English
Sentence
Interpreter
AAS-UI3a
Story
Understanding
Module
question
answer
AAS-UI3b
Question/Answer
System
AAS-UI3c
English
Sentence
Generator

AAS No.	Function	Status
AAS-RG0	Goal-driven, depth-first inference	done
AAS-RG1	Goal-drivenm procedural inference	done
AAS-RE0	Event-driven inference	done
AAS-RH0	High-level inference with meta-level knowledge	done
AAS-RF0	Function/structure-based inference	planned
AAS-RTM0	Reasoning based on instantiation	planned
AAS-RH1	Combined meta-level, function/artucture-based inference	planned
AAS-CG0	Top-down, depth-first search	done
AAS-CG1	Top-down, breadth-first search	done
AAS-CH0	Heuristic search-1	done
AAS-CH1	Heuristic search-2	design
AAS-XG0	Basic explanation (why, how) in goal-driven inference	done
AAS-XE0	Basic explanation in event-driven inference	done
AAS-XG1	Goal-driven, level sensitive explanation	testing
AAS-XC0	Explanation based on causality analysis	planned
AAS-XP0	Plan based explanation	design
AAS-KA0	Unverified knowledge acquisition	design
AAS-KA1	Simple, verified acquisition	planned
AAS-KA2	Acquisition by induction learning	planned
AAS-UI0	Fixed syntax command language interface	done
AAS-UI1	Template natural language interface	done
AAS-UI2	Case frame instantiation interface	design
AAS-UI3	Expectation-driven English parser	impl'mt
AAS-N232	RS232 message interface	design
AAS-N802.4	Ethernet (IEEE 802.4) message interface	design
AAS-OSIx	Generalized OSI interface	planned
AAS-PU0	Message/signal interface for UNIX-based systems	planned
AAS-PV0	Message/signal interface for VMS-based systems	planned
AAS-PI0	Message interface to 8086/8088 assembler module	planned
AAS-EI0	Message input/output interface to the environment	planned
AAS-EI1	Signal input/output interface to the environment	design
AAS-ES0	Input signal simulation (on/off, manual, random)	done
AAS-ES1	I/O signal simulation (on/off, manual, random, dist, v)	planned

Input Text

((David went to the store yesterday)(He got the cable)(Then he went to the Lab))

CD Graph

```
  ((PTRANS (actor (person (name (David)))) (object (person (name (David)))) (tc
(store))) (ATRANS (actor (person (name (David)))) (object (order)) (from (persc
, (name (David)))) (to (clerk))) (ATRANS (actor (clerk)) (object (cable)) (from
:clerk)) (to (person (name (David))))) (ATRANS (actor (person (name (David)))) (
ıbject (credit)) (from (person (name (David)))) (to (store))) (PTRANS (actor (ps
·son (name (David)))) (object (person (name (David)))) (from (store)) (to (Lab))
)
```

Question/Answer Session

(Did he go to the store?)

Yes, David went to a store.

(Did he get a cable?)

Yes, A clerk gave David a cable.

(Did David get a connecter?)

No, The person did not give David a connecter?)

(Did David pay money?)

No, David did not give money.

(Did David give credit?)

Yes, David gave a store a credit.