Open Research Online

Expert Systems - From Promises to Prospects: Investigations of Occupational Problem Solving Towards a More Certain Commercial Role for Knowledge-Based Systems

Thesis

For guidance on citations see FAQs.

oro.open.ac.uk

CHARLES WALTER CHURCH BA(HONS)

EXPERT SYSTEMS – FROM PROMISES TO PROSPECTS

INVESTIGATIONS OF OCCUPATIONAL PROBLEM SOLVING TOWARDS A

MORE CERTAIN COMMERCIAL ROLE FOR KNOWLEDGE-BASED SYSTEMS

PhD

AUGUST 1993

Date of submission: 16ᵗʰ August 1993
Date of award: 16ᵗʰ January 1994

ProQuest Number: 27701228

ProQuest 27701228

*With love to Valerie, Robin, Jeremy
and Tabatha the Cat.*

EXPERT SYSTEMS – FROM PROMISES TO PROSPECTS: INVESTIGATIONS OF OCCUPATIONAL PROBLEM SOLVING TOWARDS A MORE CERTAIN COMMERCIAL ROLE FOR KNOWLEDGE-BASED SYSTEMS.

CONTENTS

ABSTRACT

A range of occupational problem solvers are investigated using knowledge acquisition techniques suitable for the commercial environment. The objectives are both to understand the mechanisms by which experts routinely solve tough problems which comprise their jobs and to determine what commercially successful role expert and knowledge-based systems may take in supporting them.

A series of problem solving components are found, common between different experts, which seem to be related to the ways in which experts make problems more manageable by dividing them and solving sub-problems on a serial basis. Strategies for problem division and combining the sub-problem results are found. In some cases the nature of the problem does not permit the experts to make a good solution even when using such strategies.

Based on the investigations a range of factors are found which predispose to success or failure to implement commercially viable knowledge-based systems. Additional new potentially worthwhile areas for knowledge-based technology are suggested.

A passing commentary is made on the appropriateness of the KADS methodology for knowledge-based systems and as a component of a combined knowledge-based and conventional method.

A concluding summary is made which suggests that expert and knowledge-based systems have a good commercial future provided that the technology is carefully targeted only at those applications where it can be successful.

CHAPTER 1 - INTRODUCTION

This thesis documents a journey. The author set out on this research at a time when the champions of expert systems were predicting that they would play a major role in commercial problem solving. There was great enthusiasm fuelled by the Japanese commitment to their Fifth Generation Project and by the popular reports and predictions of authors and journalists.

The author's journey was fuelled partly by his academic interest in the behaviour of occupational experts and partly by his ability to engage a successful career in the new world of commercial expert systems. The objectives of this research have remained largely unchanged throughout its course. These were to investigate the nature of the problem solving behaviour of occupational experts and to determine how such behaviour could be best investigated, using appropriate techniques of knowledge acquisition, as a route to providing commercially useful expert systems.

Although the objectives have remained unchanged the scope for investigation widened very considerably as a result of employment opportunities. Initially, and when an employee at British Telecom, the author concentrated his efforts on investigating and replicating the skills of electronics diagnosticians. The work was successful and, with changes of employment first to management consultancy with Coopers & Lybrand and then to systems development at Hitachi, the author investigated many more occupational tasks and potential applications for expert systems. A worthwhile number of these investigations came to fruition and the resultant

expert systems were built. However, for various reasons, many of the investigations were either curtailed for lack of sponsorship or because no feasible system could be implemented.

And the author's story paralleled the changing fortunes of expert systems in the wider world. Major public investment from the Alvey and ESPRIT initiatives led to the development of a considerable number of widely reported expert systems. Good tools, methods and experience emerged from these initiatives and the explosion of expert systems usage seemed just around the corner.

But it has not happened. The world of expert systems is not dead. However, it has failed to produce the vibrant future that was so strongly predicted. Of course there are good examples of expert systems in active and profitable use. Of course there has been the absorption of much expert systems technology into mainstream computing. Of course there are still companies making a good living out of expert systems products and consultancy. But the grand promise has stabilised into a respectable niche market.

So what went wrong? Certainly there was no lack of excellence in the majority of the research and the commercial products. Some of the activity still continues in, for instance, the work on the KADS, the acronym for a Knowledge, Analysis and Documentation System, (Schreiber, 1992) methodology which offers a real route to the systematic building of expert systems. And the tools were good. KEATS (Motta, Rajan & Eisenstadt, 1988), Shelley (Anjewierden, Wielemaker et al, 1992) and the rest. But where were they used? Of course, good examples will be quoted by their sponsors. But none of the methodologies or tools have achieved anything like a worthwhile

level of commercial usage. Most of the relatively few expert systems in current use have been the work of dedicated knowledge engineers using relatively primitive tools and techniques.

And what is the author's position? He has spent the last few years of his career trying to make expert systems work in practice. And like the majority of the commercial expert systems community, he has never used KADS or any of the smart toolkits to assist his work. Much as he admired the ongoing work on these tools or projects, none of his projects seemed "quite right" for their use.

So what has this research discovered? A score of possible expert systems opportunities are reported in what follows. In some cases just the bare bones of the problem have been investigated, in others a full expert system has been successfully implemented and deployed. And there are a range of intermediate cases. And the findings fall into several categories.

First, some relatively mundane findings on how to acquire and organise expertise under conditions of commercial adversity and pressure. Not all situations are like the academic laboratory and it is often necessary to undertake at least the initial stages of knowledge acquisition on a rapid basis just so as to enable a case for funding to be put to the system sponsor.

Second, some parallels began to emerge in the way that experts organise themselves to solve routine commercial problems. And, perhaps, to cope with the limitations in their own cognitive abilities. Investigations of different problem solvers began to show common "components" in their problem solving architectures. For instance, there were components which they used

to constrain a problem to make it manageable. And, once constrained, how to progress various aspects of the problem separately and then, where appropriate, how to combine the results into an overall conclusion.

Third, a whole range of issues and obstacles evolved which were critical for the success or failure in individual expert systems. And, again, a pattern began to emerge and the issues started to be repeated in totally different expert systems opportunities. In some cases overcoming the obstacle led to a successful implementation. In others it was the key factor that prevented the system moving to a full implementation, or even being fully investigated.

The pattern of this thesis is straightforward. Chapter 2 briefly sets out the background to this research and points to the dilemma outlined above. The tools and methods seem good but they are hardly used and there are relatively few expert systems in active commercial use.

Chapter 3 is a long chapter which reports on a considerable number of expert systems investigations for which the author has been responsible. A standard format is used and the tasks are reported roughly in order of increasing complexity. During these reports there is a progressive identification of the problem solving components that recur again and again. And, as part of the commentary on each investigation, note is taken on the key issues, overcome or otherwise, which were crucial to the success or failure to implement an expert system based on the investigation.

Chapters 4, 5 and 6 take up separate aspects of the findings resulting from the investigations reported in Chapter 3. Chapter 4 is a record of the practical experience and techniques used and developed in order to undertake the investigations. Much of it is concerned with knowledge

acquisition in the commercial situation. The author has been very fortunate in that his various employments have provided him with a rich variety of expert systems opportunities hardly likely to be encountered by the full-time academic researcher. However, the realities of the commercial situation often mean that experts are difficult to access and that a very rapid "scoping" of the knowledge has to be undertaken before a full system will be funded. Chapter 4 reports on the tricks and compromises that the author has found necessary to achieve his objectives under situations of commercial pressure.

Chapter 5 is a relatively short chapter summarising the "problem solving components" progressively identified in Chapter 3. What strategic combinations of such components do experts use to overcome the limitations in their own abilities? What is the commonality of these components between experts in widely different fields? Is there any value in recognising such common components when implementing expert systems? Or are they just artifacts of the human experts better left behind when the expertise is transformed for the machine? The findings presented in Chapter 5 are relatively clear insofar as the nature and purpose of the problem solving components found in Chapter 3. But there are many unanswered questions as to the value of the findings for the commercial expert systems implementor or theorist.

Chapter 6 presents the third and final perspective on the findings of Chapter 3. A range of factors, in common topic groupings, are presented which either predisposed to feasibility, success or failure of particular expert systems implementations. Where success was achieved, or could have clearly been achieved with the sponsor's goodwill, then the enabling technique or method is clearly identified. Conversely, where a major

obstacle prevented a commercially feasible implementation then this is described together with suggestions, if any, of ways which might overcome similar obstacles in future.

A further particular topic is considered as part of Chapter 6. The investigations in Chapter 3 have highlighted a number of opportunities where expert systems need to rest on, and communicate between, the knowledge of two organisationally separate experts. A key example is that provided by the designer and diagnostician of a piece of electronic equipment who have the same domain of interest (ie the particular circuit) but who never meet and communicate poorly to their mutual working detriment.

Chapter 7 takes an optimistic stance for the future and summarises a series of relatively straightforward measures that could recover some of the promise of expert systems. Actions in support of some of these measures are already in progress and it may be that more widescale usage of expert systems, maybe by another name, will be seen not too far in the distant future.

Chapter 8 is a short concluding summary of the work and is followed by Acknowledgements and a Bibliography.

CHAPTER 2 — BACKGROUND TO THE RESEARCH

2.1   INTRODUCTION

This chapter sets the scene for the remainder of this thesis. First, there is a review of a spectrum of academic and commercial interests ranging from human simulation to knowledge based systems. On the one hand there are those, mainly from academia, who are motivated to understand the workings of human cognition in the field of problem solving. On the other hand there are those who want to use knowledge acquired from humans as the basis of software to provide business advantage. The author stands in both camps.

Second, there is a brief review of some current theories and models of expertise. This leads directly to a discussion of KADS and supporting toolkits for knowledge engineers. There has been much research and applied progress in recent years and expert and knowledge-based systems should by now be deeply entrenched in commercial use.

But, in the brief third part of this chapter, we note that the promise has not been delivered. Expert systems still have a tentative role in commercial information technology and there is little sign of the long-expected explosion of usage.

2.2   AI, HUMAN REPLICATION, EXPERT SYSTEMS AND KBS

There is reasonable consensus on a core definition of Artificial Intelligence (AI) in a form typically such as "the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the

characteristics we associate with intelligence in human behaviour – understanding language, learning, reasoning, problem solving, and so on." (Barr & Feigenbaum, 1981). Nevertheless, although there is agreement on the nature of AI, there have always been considerable differences in the motivations of those pursuing the study of AI. For some, including the author, AI is seen as a means of exploring and testing ideas on the nature of human cognition. In our own case we are interested to discover how humans solve (or make errors in solving) real life problems.

The "human replication" sub-set of AI is concerned with the development of computer programs that not only deliver the same results (ie surface appearance) of intelligent humans but also have underlying mechanisms analogous to those of the human problem solver.

In the rest of AI there is no such attempt at human replication and, quite often, there is forthright determination to "better" the human process (eg Nilsson, 1971). Commercial sponsors are generally not interested on research into human cognition. They primarily see AI as a powerful way of attacking business problems not previously amenable to "computerisation". Specifically "expert systems" (considered fully in texts such as Hayes-Roth, Waterman & Lenat, 1983) can be considered as computer programs, within established AI technology, that simulate or support the skill of a single occupational human problem solver.

As we understand it expert systems:

(a)   are rarely concerned with human replication; competent performance is the sole criterion and representations and mechanisms analogous to the human are not required.

8

(b)   are restricted to derivatives of the problem solving strand of AI (ie the strand originating from game players, journey planners etc). Expert systems are not concerned with perception, natural language understanding, robotics etc although these other AI sub-fields are sometimes involved in an integrated expert systems implementation.

(c)   are motivated by commercial/business needs and applications. As such only feasible applications within the currently established AI technology are normally considered. For instance, expert systems based on neural computing will become commercially viable only if and when that technology becomes sufficiently generalised and reliable for widespread commercial use.

(d)   are intended to simulate (or support) the experience of a single human problem solver. This is a crucial limitation and implies that an expert system is "sufficient" if it matches the competence of the expert on whose knowledge it was based. It also implies that expert systems, although they may be based on the expertise of more than one expert, do not necessarily simulate the co-operative problem solving skills of a community of experts.

The term "knowledge-based systems" (KBS) is wider than the term "expert systems" and covers systems which employ for their processing knowledge acquired from humans but which have both an external appearance and an internal architecture and mechanisms that may have nothing in common with any human expert. Jackson (1990) provides a useful example to distinguish

between the generality of a KBS and the specific behaviour of an expert system. He cites a program capable of conversing about the weather as being a KBS, providing it performs its task by applying heuristics to a symbolic representation of knowledge, even if it did not embody any expertise in meteorology. However, an expert system in the domain of meteorology ought to be able to provide weather forecasts or some other output equivalent to that of a human expert.

Thus, in outline summary, human replication demands both performance and, insofar as is possible, internal mechanisms analogous to that of the expert. Expert systems require only elements of external performance equivalent to that of an expert and knowledge-based systems gainfully use knowledge acquired from humans as software building blocks to deliver performance that may or may not be anything like that of a human expert.

2.3 CURRENT THEORIES AND MODELS OF EXPERTISE

Theories on the nature of expertise, and how it should be modelled, are seen by the author as being made in two orthogonal but complimentary directions. The starting point for both progressions is the all-purpose generalist representation of knowledge which has long pervaded the expert systems world. Newell (1982) started the break with this tradition with "The Knowledge Level" where he emphasised the need to describe the types of knowledge required to perform a reasoning task. The first direction, taking an approach involving Generic Task Knowledge is described first and is followed by a brief review of those advocating the Structured Roles of Knowledge. This latter topic leads into a discussion of the KADS methodology which is now of increasing importance for the principled

builder of knowledge-based systems. A further discussion relevant to KADS is included in Chapter 7 where we review the implications of the research described in this thesis.

Generic Task Knowledge

In one direction Chandrasekaran (1987) outlines a Generic Task approach which is claimed to provide a better framework for building expert systems. He claims that the level of abstraction currently used for building knowledge-based systems is too low to provide a sufficiently rich vocabulary for knowledge and control. The current level, with rules, frames, logic etc is something of an assembly language for knowledge.

What Chandrasekaran claims is needed are higher level building blocks each appropriate to a particular type of problem solving. Such building blocks each map onto one of a number of "generic tasks", such as diagnosis, and their use both provides a richer and more appropriate vocabulary and also offers a more natural route for knowledge acquisition.

In Chandrasekaran's scheme of things each representation has a natural family of inference mechanisms that can act on it. Together they form an architecture and, if the inference is fixed, the resulting inference scheme forms a generic shell into which the appropriate knowledge can be inserted.

Chandrasekaran provides a first-rate report of the way in which real expert systems knowledge bases do not explicitly reveal the control structures which the knowledge engineer has implicitly pre-fabricated to exploit the peculiarities of, for instance, the backward or forward chaining mechanisms

11

provided by a shell. He also points to the considerable pre-occupation with conflict-resolution strategies to manipulate control. Such contrived control would not be necessary if the knowledge were properly organised (ie modularised) to tight relevant bodies of knowledge.

Chandrasekaran provides a detailed introduction to Generic Tasks and, in particular, the Generic Tasks involved in diagnostic reasoning. In passing he mentions the "Fallacy of Surface Phenomenalism" which results from the dumb simulation of expert behaviour without any awareness of the underlying internal architectures of the generic tasks simulated.

Chandrasekaran's work seems to be similarly motivated and in the same general direction as Clancey's (1985) work on Heuristic Classification. However, although Chandrasekaran focuses more-or-less exclusively on diagnostic reasoning as a prime example for the Generic Task approach, he has a far wider perspective than Clancey.

Clancey also approaches the situation in reverse. Whereas Chandrasekaran is unable to accept the inadequacies of the "assembly language" rules etc and sees a range of Generic Tasks with corresponding architectures, Clancey (albeit with great elegance and detail) merely notices that a class of programs happen to all reduce to what Chandrasekaran may have called a single Generic Task Architectures.

In summary, Chandrasekaran, Clancey and others are moving the all-purpose generalist representation of knowledge in the direction of task-type-specific shells and architectures.

Structure Roles of Knowledge

The other dimension of movement is spearheaded by workers such as Wielinga and Breuker (1986) who emphasise the different structural roles that knowledge can play in the reasoning process. Wielinga and Breuker see a four-layer structure of knowledge. The layers represent domain knowledge, inference knowledge, task knowledge and strategic knowledge. Other writers, for example Steels (1988) and Pople (1982), have different emphases and use different terminology but basically all propose a differentiation of knowledge according to its roles in the reasoning process.

Wielinga and Breuker make a major positive statement in their demand that any descriptional framework for knowledge should be at the epistemological level. Personally we do not have any objection to knowledge engineers when they make the occasional colloquial description of humans as "forward chaining" etc. Such terms now have an informal meaning over and above the implementational meaning.

Nevertheless, at the more serious level, Wielinga and Breuker's criticism of current AI descriptions which repeatedly gravitate to the implementation level are extremely well-founded. KADS, which rests heavily on their work, is discussed immediately below.

2.4   KADS KBS DEVELOPMENT METHODOLOGY

KADS (Wielinga, Schreiber & Breuker, 1992) is a structured methodology for the development of knowledge-based systems. It is the result of a very effective collaboration between strong academic and industrial partners. The

13

success and promise of the work has led to major ESPRIT and other funding. In addition to its existing commercial users, KADS is now under active consideration by many developers of advanced software systems.

KADS is based on the premises that knowledge engineering is a difficult task and that a structured route must be provided to transform human expertise into runnable knowledge-based systems.

KADS uses a modelling approach to cope with the complexities of the existing human and organisational situation and its successive transformation to a situation employing a knowledge-based system.

Modelling is useful since it allows certain aspects of a target knowledge-based system to be emphasised whilst, at the same time, ignoring other aspects which may be committed to another model. (Wielinga, Schreiber & Breuker, 1992) claim that the use of multiple models allows the knowledge engineer to use a "divide-and-conquer" strategy to cope with the complexities of the knowledge engineering process.) Thus, for instance, an Organizational Model has only to consider the socio-organizational environment in which the knowledge-based system will be required to perform its task. In contrast, the Application Model describes what problems the target application should solve in the organization, what should be its organizational function and what are its organizational constraints in terms of speed, hardware requirements and integration with other systems. In some respects, the Application Model is the closest thing to a "Job Description" for the knowledge-based system.

The Application Model is related to the Task Model in that the Task Model demonstrates how the functions specified in the Application Model are

14

satisfied by the tasks performed by the system.

Tasks are usually decomposed into sub-tasks and primitive tasks that are themselves modelled by the Model of Cooperation and the Model of Expertise. Here again the use of different models, covering the separate aspects of "cooperation" and "expertise", offers the knowledge engineer powerful assistance in understanding and manipulating the content of the evolving knowledge-based system.

The Model of Cooperation provides a specification of those tasks where cooperative effort is required. For instance, a knowledge-based electronics diagnostic system must rely on a human user to actually perform components substitution. Similarly, a financial trader may want certain Reuters information to be "volunteered" to a knowledge-based system since this is exactly the sort of information to which he would be continually alert.

The construction of the Model of Expertise is the core enabling activity in knowledge-based system building. It is also the key distinguishing feature that separates knowledge-based system building from conventional systems development. We will consider the Model of Expertise in a little more detail since some of the work reported in this thesis considers knowledge which is outside the Model of Expertise.

The Model of Expertise is expressed in terms of the behaviour that the target system should display and the knowledge that facilitates the generation of such behaviour. The specification of the Model of Expertise is at the "knowledge level" (Newell, 1982) rather than the symbol level where the emphasis is on computational techniques and representations, such as rules and frames, that will generate the behaviour of the implemented

15

system.

A key aspect for our consideration is that the Model of Expertise and, for that matter all of the other KADS models, are biased toward what the target knowledge-based system should do rather than how a particular human expert behaves. KADS is a methodology whose objective is to produce viable knowledge-based systems and it would be unusual if it took a backward look in the form of a Cognitive Model of the existing human expert. Hence the use and appropriateness of the term "knowledge-based systems" rather than the term "expert systems" which may have a connotation of modelling a particular expert.

Formulation of the Model of Expertise (and, more in theory than practice, other models) is much assisted by the availability of libraries of re-usable model elements. Two particular types of re-usable elements are available in KADS. The first are primitive problem solving actions such as the instantiation of a concept to an instance. In KADS such primitives are termed "knowledge sources". (We use this term and define it in Chapter 3 with a wider meaning.)

The second type of re-usable element is larger than the knowledge source and is termed an interpretation model. Interpretation models are partial models of expertise which lack the details of actual domain models. A decision tree referencing a taxonomy of task types is used by the knowledge engineer to select an appropriate interpretation model from the library. Once an interpretation model has been selected it may be used directly or in a refined ("differentiated") form to act as a host for the acquired domain knowledge.

In summary, once an interpretation model has been correctly selected it is used as a top-down receptacle and the knowledge engineer will seek the relevant domain knowledge to populate it. Obviously this method of working has great potential advantages for the knowledge engineer who can systematically acquire the domain knowledge for the selected model rather than attempt to make sense of and organise a mass of randomly acquired domain expertise.

The Model of Expertise and the Model of Cooperation are combined to form a Conceptual Model of expertise and cooperation expressed in a implementation-independent language.

The Design Model takes the leap of expressing the specified behaviour in terms of the target computational framework. The separation of the Conceptual Model from the realities and constraints of the implementation allows the knowledge engineer considerable freedom in expressing the full range of behaviour required from the target knowledge-based system. However, in moving from the Conceptual Model to the Design Model the knowledge engineer has to face up to the practicalities of making an implementation. Fortunately, KADS provides support for this process, in the form of notations etc, in a way which preserves the structure of the Conceptual Model through to the Design Model. Once the Design Model is completed it is a comparatively straightforward process to proceed to the final knowledge-based system implementation.

KBS does not easily match the expectations of the users of conventional systems methods. Saward, Land & Bingham (1993) report that is difficult to test systems feasibility without knowing the nature of the knowledge in advance, definition of systems requirements is difficult and knowledge

acquisition may discover knowledge or information which, when analyzed, leads to a mid-project revision of the system design. The CommonKADS life-cycle model (de Hoog, Martil, Wielinga, Taylor, Bright & van de Velde, 1992) is currently leading to a range of products that are intended to address such problems.

Methodologies such as KADS are usefully supported by computer-based tools. These are the equivalent to the CASE (Computer Aided Software Engineering) tools used to support conventional systems analysis and design methods. In the case of KADS, the Shelley workbench supports the entire modelling process including the selection and instantiation of an interpretation model to form a Design Model. Shelley and other support tools for knowledge engineering are briefly discussed in the following section.

2.6 COMPUTER-BASED TOOLS FOR KNOWLEDGE ENGINEERS

The use of computer-based tools is not discussed in this thesis. However, some of the implementations resulting from the investigations described in Chapter 3 were made with the aid of computer-based tools. Such tools are briefly discussed here as a review of an important topic closely related to this research.

Probably the commonest methods and tools are those associated with the major commercial expert systems shells. In the main such methods and tools are equivalent to "lower CASE" tools and support the implementation and maintenance of knowledge bases. Examples are provided by KEE, ART, KnowledgeCraft, KAPPA, Nexpert and ObjectIQ.

Tools such as KEATS (Motta, Rajan & Eisenstadt, 1988) and Shelley (Bouchet, 1989), the KADS workbench, attempt to cover the wider knowledge engineering process by offering knowledge acquisition tools and methods of understanding and manipulating various models of knowledge.

KEATS has a particular mandate to support the knowledge engineer undertaking the implementation of larger expert systems. As such it offers particular forms of what we classify as "problem solving props" in Chapter 4 to structure the knowledge engineer's work. By offering tools which allow work on the organisation of selected sub-domains of knowledge it helps the expert knowledge engineer to cope with the complexities of the domain knowledge.

The tools so far discussed are essentially passive and neutral in their support. In other words they provide a supportive and hygienic medium for the knowledge engineer to undertake his task. Other tools take a more proactive or biased approach and attempt to offer "domain-specific" or "principled" guidance to the knowledge engineer.

Thus in the first "domain-specific" category, tools such as ROGET (Bennet, 1983), MOLE (Eshelman, Ehret, McDermott & Tan, 1988) and OPAL (Musen, Fagan, Combs and Shortliffe, 1988) have structures and expectations in respect of knowledge concerned with a particular domain, such as diagnosis, or style of problem solving.

Similarly, in the "principled" category there is TEIRESIAS (Davis, 1979), which has a model-driven expectation of the rules required for an effective implementation.

Further "intelligent" tools, offering co-operative support, attempt to partially take over the work of the knowledge engineer. For example, although very limited, PROTOS (Bareiss & Weir, 1988) has sufficient knowledge to attempt a principled construction of the knowledge required to support a novel case. If the system fails, or partially fails, the user or knowledge engineer must provide PROTOS with the missing knowledge.

2.7   COMMERCIAL UPTAKE OF EXPERT AND KNOWLEDGE-BASED SYSTEMS

From the above it would seem that everything is in place for the widespread adoption of expert and knowledge-based systems by commerce and industry. The more principled tools, such as those described immediately above, have long been supplemented by a welter of commercial expert systems software ranging from simple PC rule-chaining shells to sophisticated "AI" toolkits. Additionally there has been major promotion and funding of knowledge-based initiatives from UK government departments including the dedicated Alvey Directorate. Further money has been available from the EC which has sponsored a series of international collaborations between academic and industrial partners as part of the ESPRIT programme.

Yet the commercial uptake of expert and knowledge-based systems is very disappointing. Many glowing reports originated from the Alvey "major demonstrators", there was a proliferation of special interest clubs and there are often repeated stories of commercially successful expert systems in companies such as American Express etc. Many expert systems have indeed been commercially deployed, a small proportion of which are a result of the author's own commercial work. Yet the overall level of expert systems adoption is dismal.

Because of his commercial work the author is still fortunate enough to be able to penetrate the information technology departments of many top corporates in the financial and manufacturing sectors. In some cases there remain small groups of talented expert systems workers who have limited success at promoting the technology in the mainstream business departments. In others companies such workers no longer refer to themselves as "knowledge engineers" and are chasing new goals in object-oriented technology and similar areas.

And it is the same for the vendors of AI/expert systems toolkits. There have been many bankruptcies and most of the vendors who remain have repositioned themselves away from expert systems and claim more general utility for their renamed products as advanced toolkits for general software development.

So why is this? The research and development has been well-founded and the results have been very adequately promoted and financed in commercial situations. The business potential of applying knowledge in expert and knowledge-based systems is well understood and there has been surprising goodwill and finance from industry in unprecedented industrial-academic partnerships.

The author has been an active participant in the expert systems community and remains confident that the technology can offer real advantages to business. In the chapters that follow he documents his experience, partly research, partly commercial, in order to highlight some of the factors that can convert the promise of expert systems to prospects of widescale adoption.

CHAPTER 3 – INVESTIGATIONS OF OCCUPATIONAL PROBLEM SOLVING

3.1   INTRODUCTION

This chapter reports on fourteen problem solving tasks from a wide range of occupational environments. The author was the primary instigator of the investigation of all the reported tasks although, as explained below, not all progressed to full expert systems and, for those that did, the author usually managed the implementation rather than undertaking it himself.

Opportunistic Nature of Task Investigations

The applied problem solving tasks considered here arose on a semi-opportunistic basis as a result of the author's career experience. Considerable care has been taken to exclude those matters which are in commercial confidence or which were wholly or partly the work of others.

Many of the tasks actually resulted in expert systems implementations and such commercial outcomes are reported. However, some investigations could not or did not progress to implementation but, nevertheless, led to useful insights into the nature of the problem solving involved or to the possibilities of implementation. Implementation outcomes, including the reasons for failure to implement, will be used as the basis for the discussions on expert systems implementation feasibility in Chapter 6.

Two Special Personal Projects

The first tasks to be investigated (Telecom Repair Service and Benchtop Electronics Diagnosis) were the result of selection by the author from tasks available for study as a result of his then employment with British Telecom and are given extended treatment in the following sections. In both cases a PC-based expert system shell arose out of the work in addition to the particular expert systems. In the case of Telecom Repair Service this was the "DIS" shell for which the author was wholly responsible. For Benchtop Electronic Diagnosis it was the "TRACKER" shell for which the author was joint designer but not implementor of the commercial version. These shells were developed at a time when commercial expert systems shells were still under development and there is little to report about them that is now not available in commercial products.

Exploratory and Managerial Role for Other Tasks

Apart from Telecom Repair Service and Benchtop Electronic Diagnosis, in most of the other cases the author was the manager of a commercial expert systems group and his involvement was usually twofold. First, there was a responsibility to instigate and then manage the investigation of the tasks and to carry this forward to an expert system implementation in those cases where this could be achieved or sponsored. In most of these cases the author had only a partial personal role in the detailed later stages of the investigation and implementation. Such details are not reported for this and commercial confidentiality reasons.

The second part of his involvement in these later tasks is far more relevant

24

to this research. In fact it has given him an investigative opportunity which might not have been available if he had been involved fully in the later detailed extended work of any of the individual implementations.

In all the reported tasks the author has been responsible for the initial knowledge acquisition and knowledge analysis for the tasks. In many cases this was associated with his professional responsibility for assessing the feasibility of a commercial expert system implementation and "scoping and estimating" for the cost of the work.

Tasks Investigated

The following were the tasks investigated and reported in this thesis. Although they are in the main generic titles of tasks found in many organisations, the titles below will be used throughout this thesis solely to describe the specific tasks investigated.

COMMERCIAL BUSINESS UNDERWRITING (Insurance Company)

FOREIGN EXCHANGE EXPOSURE (Aero Engine Manufacturer)

BENCHTOP ELECTRONIC DIAGNOSIS (British Telecom)

AUDIT WORK PLANNING (Coopers & Lybrand)

BUSINESS PERFORMANCE ASSESSMENT (Government Health Department)

MULTINATIONAL TAX PLANNING (Chemical Manufacturer)

25

ENGINEERING PREVENTATIVE MAINTENANCE (Cigarette Manufacturer)

PERMANENT HEALTH INSURANCE (Insurance Company)

LIFE UNDERWRITING (Insurance Company)

SMALL BUSINESS GUIDANCE (Government Training Agency)

PENSIONS MANAGEMENT (Public Utility Pension Scheme)

JOB SHOP SCHEDULING (Electronic Consumer Products Manufacturer)

TELECOM REPAIR SERVICE (British Telecom)

EUROBOND WARRANT TRADING (International Bank)

Non-Trivial Occupational Tasks and Problems

All of the fourteen tasks and problems considered in this chapter are non-trivial examples of current occupational problem solving. In each case either a task is described which represents the primary activity for which the task-holder is employed or a problem is highlighted which has major business and financial implications for the company concerned. In the latter case, the problems are those for which the company has only a partial solution and for which, if it were technically and economically feasible, it would have been advantageous to give expert systems support.

Although some of them would find the term a little strange, all the

individuals concerned with the tasks can be considered as "experts" and this is the term we will use for them. In other words, they are all custodians of knowledge and experience, either learned or self-taught, which they apply to solve problems in a way far excelling non-experts.

Commercial Confidentiality and Professional Skills Issues

Commercial knowledge is the subject of this research and it is also a major asset (perhaps the primary asset) of the organisations where it is exercised. Accordingly the author has had to use considerable care to focus only on matters which are reportable in a public document such as this thesis.

Special conditions applied for Telecom Repair Service and Benchtop Electronic Diagnosis. As a result the knowledge associated with these tasks is reported (or, more properly, sampled) in detail.

In other cases detailed commercial knowledge is not reported except in a rather fragmentary way. Instead, and in line with the primary focus of this research, only the strategic and generic content of the task knowledge is reported here. Such strategic and generic knowledge is in the public domain insofar as these tasks are concerned.

The division between the detailed commercially confidential knowledge and the personal professional skills of the experts is very relevant to the above issues of commercial confidentiality. As mentioned above, we have now investigated some "replicas" of the tasks reported in this research. We have also discussed the nature of some of the reported problem solving tasks with experts in organisations separated from the source of the

reported expertise. In all cases it seems that the style of the task problem solving which we have discovered is very largely similar in the different, and often competitor, organisations.

In fact it seems that, almost without exception, the tasks we have investigated are better classified as ""professional skills" ("owned" by the experts) than "proprietary practices" ("owned" by their employing organisations). In other words they are the sort of skills that the relevant experts would usefully carry with them from employer to employer. In fact it is such skills and task-generic knowledge that make experts employable in a market economy job market. The fact that their "knowledge CVs" have been accumulated as a result of specific confidential experiences with a particular company does not prohibit their using the generic product of such experience with their next and future employers.

For instance, our report on Life Underwriting embodies the sort of knowledge that would be expected for an applicant for a professional life underwriting post rather than the task specific knowledge of "Insurance Company X". Similarly the reported material on Pensions Management is that which would be recognisable for any competent applicant for a post of "Pensions Manager". Although derived from a specific task in a particular commercial organisation our report is not specific to "Company Pensions Scheme X".

Thus, it seems that the style (ie primarily generic and strategic knowledge) of many occupational tasks is common (and therefore in the "public domain") across those organisations where the task is practised. What is commercially confidential is the specific detailed knowledge and, in particular, the values instantiated in the knowledge.

Examples of "Public" and "Commercial" Knowledge

The following examples should highlight the two aspects discussed above. First, they should illustrate the principles both on which we have decided that which is reportable in this thesis. Second, and more important, they should illustrate what would separate that knowledge either "built in" to an application shell or remaining to be implemented for a specific commercial expert system implementation.

As a first example, in our account of Commercial Business Underwriting we report that "the underwriter's basis of quotation contained categorisation of company type (eg sporting event caterers) and rates from risk tables". This would be no surprise to any competitor company. Similarly they and any other commercial business insurer would require any human expert or expert system to take account of such categorisations. It is even the type of knowledge that might be required to pass underwriting examinations. However, what would be confidential is the revelation that Insurance Company "X"'s commercial business Employee Liability rates for sporting event caterers were calculated at 23% above normal rates for London and the South-East and at 18% above normal elsewhere.

As a second short example we report in our account of Pensions Management that "a child may be a dependent student" but do not give the specific information on which full-time or part-time courses qualify for studentship (eg in excess of twenty hours per week leading to a recognised academic qualification above A-Level standard). This is commercial task-specific information but the principle, that a dependent student may be deemed a "child", is public domain generic knowledge.

Format of Task Reports

The reports on Telecom Repair Service and Benchtop Electronic Diagnosis are reported in some considerable detail. However, for the majority of tasks, where a system was implemented or a worthwhile investigation was made, a standard report format is used for purposes of comparison between tasks. In the remaining few tasks, where only a superficial investigation was made, then the standard format is used in an abbreviated way. The topics covered in the standard format are:

Introduction - general introduction to the investigation and the background to the task and its purpose in the organisation.

Sponsoring Organisation - a description but, generally, not identification of the organisation where the expertise is used.

Author's Role - to clarify the nature of the author's personal role in the research and to indicate any further commercial work, not reported here, by the author and his colleagues.

Outline Task Description - a description of the inputs and outputs to the task and the normal handling of the task by the expert.

Form of Investigation - the extent of the research investigation and the manner in which it was undertaken.

Findings on Nature of Problem Solving

General - a short overall description of the way in which the

30

expert handles the task problem solving.

Problem Structuring Props (Tools and Aids) – relating to the discussion on Problem Solving Props at the end of Chapter 4.

Generic Data – those sources of data which the expert uses and which are not specific to any particular task item under consideration.

Specific Data – the specific data, often partially comprising the problem description, used in the solution of a particular task item.

Knowledge Sources – this term is used in a special way as explained below in the section on "Expert systems explanation of problem solving".

Open/Closed Problem – indicating whether or not there. is a finite set of problem descriptions and/or solutions even if this set is very large.

Proactive/Reactive Problem – to indicate if the problem solving is initiated by the presentation of specific problems to the expert or, far less usually, by some overriding ongoing business goal.

Problem Solving Architecture – a high-level description of the problem solving architecture including identification of key problem solving

components.

Diagrammatic Representation of Problem Solving Components - the author's own diagrammatic "shorthand", developed progressively as an increasing number of tasks were investigated. The diagrams, often produced early in the task investigation, enabled him to make ready characterisation and comparison of the problem solving architectures of the various tasks. The recurrent major components of the problem solving architectures were recognised by the author as more tasks were investigated and are discussed and summarised in Chapter 5.

Commercial Implementation Details - to indicate what was the commercial result, if any, in the form of an expert systems implementation.

Additional Comments - to summarise particular points of interest.

Use of Terminology and "Expert systems" Explanations of Human Problem Solving

As discussed in the Background to the Research section, one of the origins of AI was to use computer simulations to discipline and test theories of cognitive structures and functions. One of the outcomes of AI is that it has given us a vocabulary for the computational work that is sometimes perhaps richer (or more familiar) than the vocabulary for the description of the human problem solver.

In the task descriptions that follow we have rather deliberately used AI and

expert systems terms for descriptions and explanations of human problem solving. Thus we refer to "search", "knowledge sources" and "forward reasoning" etc in contexts entirely divorced from issues of implementation.

For instance, we use the word "search" in a fairly loose and general way. Nevertheless our usage is reasonably straightforward in meaning. In this chapter, we make particular usage of the term "Knowledge Sources" as one of the standard topic headings for each task description. In this case the usage is a little different from its usage in relation to blackboard systems, structure of knowledge as in, for instance, Clancey (1983), etc. It also differs from the special meaning used in the KADS methodology and described in Chapter 2.

In our usage a Knowledge Source is a reasonably self-contained part of the expert's knowledge which, at least for part of the problem solving, can act independently on the current state of the problem. Thus in Small Business Guidance it is clear that the expert advisor to the owner-manager can, at separate times, advise his client from the standpoint of an "accountant", "personnel consultant", "distribution expert" etc. These separate fields of expertise are more distinct than just "modules" and for this reason we label them separate Knowledge Sources.

Ordering and Content of Task Report Sections of this Chapter

There are several possible orders in which the fourteen task reports could be presented. Chronological order of the investigations is one possibility as is an ordering grouping the tasks into engineering, financial, commercial etc.

In the event the author has decided to order the tasks approximately in increasing complexity of problem solving architecture. The progress of this research has led the author to recognise common architectural components in the problem solving tasks investigated. These components are introduced progressively under the "Problem Solving Architecture" section of each task report and are fully reviewed and summarised in Chapter 5.

There are considerable differences in the level of detail of the various reports. For instance, the reports on Telecom Repair Service and Benchtop Electronic Diagnosis, which both deal with major engineering diagnostics tasks investigated at British Telecom, are in far more detail than the remainder of the reports in this chapter.

The remainder of the tasks were sufficiently well investigated to enable reporting on the nature of the problem solving and the apparent problem solving architecture used by the human expert. All of these reports, except that on the diagnostic task of Engineering Preventative Maintenance, fall within the general category of financial and commercial problem solving.

Supporting Material for the Tasks Investigated – Appendices A, B & C

The three appendices to this thesis support the investigations reported in this chapter. Appendix A is purely relevant to knowledge acquisition for Telecom Repair Service and is explained fully in Section 3.14. Appendix B provides a short tabular comparative summary of the characteristics of the tasks investigated and Appendix C, which also supports the conclusions to Chapter 5, summarises the tasks as a matrix in respect of the problem solving components recognised as the tasks were progressively investigated.

3.2   COMMERCIAL BUSINESS UNDERWRITING

Sponsoring Organisation

The investigation and subsequent implementation was sponsored by a large UK insurance company.

Introduction

The underwriting of commercial business risks, such as those associated with running factories or operating building sites, is an extremely competitive field.   Unlike personal applicants for insurance, companies invariably ask for quotations from several different insurance companies at the time of annual insurance renewal.

Provision of quotations is expensive and involves senior staff, particularly where high risks are involved.   Further, less than half of the quotations are taken up and the work involved in preparation of the remainder is entirely wasted.   For this reason an expert system implementation was considered both to facilitate the work being undertaken more quickly by more junior, less expensive, staff and also to cope with the increasing number of requests for quotation.

Author's Role

The author solely undertook all of the knowledge acquisition, analysis and scoping which resulted in the core design for the system.   However, he

moved on to other work after the initial design stage and was only marginally involved in the successful implementation of the system.

Outline Task Description

The underwriting expert receives, usually by fax, a handwritten description of a commercial business risk from one of the insurance company's regional offices. Based on his underwriting knowledge and reference to tables and documentation, he prepares a quotation to send to the regional office for onward submission to the applicant.

Form of Investigation

The opportunity to investigate this task came from the IT department of the insurance company and a good source of briefing was available from an IT manger experienced in supporting commercial business underwriting. As a result, when knowledge acquisition was initiated with the nominated experienced commercial business underwriter, most of the structure of the task, the nature of the supporting documentation and the domain "buzzwords" were already known.

This was particularly useful in a domain with which the author had had no experience, even as a customer. In contrast the main elements of life insurance, as reported later in this chapter, were understood as a "consumer" and knowledge acquisition could be undertaken directly with the underwriter.

Since the task involved considerable usage of tables and documentation (and to some extent the expertise can be considered to be in the application of documented "rules") this material, and particularly its structure, was studied in depth as part of the knowledge acquisition.

Findings on Nature of Problem Solving

Commercial business underwriting is regarded as requiring great skill and expertise and there is considerable concern about lack of skilled staff to meet the demand for quotations.

However, perhaps a little sadly, on full investigation it was discovered that the task is not a particularly difficult or complex one. What seems to make the task esteemed, and therefore "expert", is that the financial risks are very great. Since the underwriter is dealing with risks of many hundreds of thousands of pounds it is automatically assumed that great judgement and knowledge are required. In fact, as stated above, most of the skill is in invoking rules in the documentation and, in fact, very little discretionary judgement is involved.

Each problem is described by the branch office's description of the risk. It is rarely necessary to seek further details to make a quotation. Where there is any ambiguity in the description (eg as to the presence of fireproof walls) the quotation is based on the assumption that a good situation applies and that the quotation is subject to the company's inspection of the premises.

INVESTIGATIONS OF OCCUPATIONAL PROBLEM SOLVING

Problem Structuring Props (Tools and Aids)

This task was unique in that it demonstrated the use of what we term a "virtual prop". The prop was embodied in the structure of the underwriter's free-hand basis of quotation. This document was not the brief quotation calculation submitted back to the branch office but contained the "working" similar to that in a piece of school mathematics.

Starting with a blank sheet of paper, the underwriter would assemble this basis of quotation in exactly the same way on every occasion. In every respect the plain paper was like a "form" where the "boxes" were only visible to the underwriter.

The underwriter's basis of quotation contained the following main fields:

1) Name of applicant company

2) Note of any special risks

3) Categorisation of company type (eg in-flight caterers) and rates from risk tables

4) Individual liability (public, employer's, etc) calculations based on factors such as turnover or number of employees

5) Premium calculation based rates and liabilities

6) Discount calculation based on size of firm

Generic Data

Generic data was in the form of the already mentioned risk tables (comprising ratings for all categories of business and liabilities) and special

rules, for instance, for discounting for large customers. The risk tables can be considered as a form of "compiled knowledge" based on the statistical analysis of claims on different types of commercial risks over many years. Such "knowledge" could never be effectively accumulated by a single underwriter but is the legacy of the risks insured by many underwriters over a great number of years.

Specific Data

Apart from the invitation to quote, the expert advised that information from previous year's quotations was available. However, although these records were kept at some considerable expense, there was no evidence that this applicant-specific data was ever used. All quotations appeared to be constructed laboriously from first principles on the basis of currently submitted information. The use of records from previous years seemed to be just part of the expert's "mystique" to which he unwittingly contributed.

Knowledge Sources

Although there are separate fields in respect of liability (eg employers liability or contract works liability) and discrete fields of knowledge on rating and discounting, the expert used all the knowledge in a fully-integrated way and thus only one knowledge source is envisioned.

Open/Closed Problem

The problem is closed in that there are finite set of business types and risks associated with them. Any novel cases are converted to analogous cases within the closed set. Thus the underwriter assigned a satellite dish installer as a "roofer" (for which a rate existed) since the inherited primary risk (ie falling) is the same.

Proactive/Reactive Problem

The problem solving is solely reactive to the invitation to quote.

Problem Solving Architecture

This task and Foreign Exchange Exposure, reported in the next section, both demonstrated the most simple form of search encountered in all of the occupational problem solving tasks investigated. Only one knowledge sub-domain is involved and this is completely searched more-or-less algorithmically. The search involved can be represented by the simplest expert system shell, using either backward or forward chaining, with discrimination tests performed both at intermediate and terminal nodes.

The task can be regarded as expert because, although inference and search space are straightforward, the size of the search space can be very large. Although the search is shallow it is also very wide as a result of many different types of business. Thus, for example, in Commercial Business Underwriting the search can be represented as shown on the next page:

40

COMMERCIAL BUSINESS UNDERWRITING
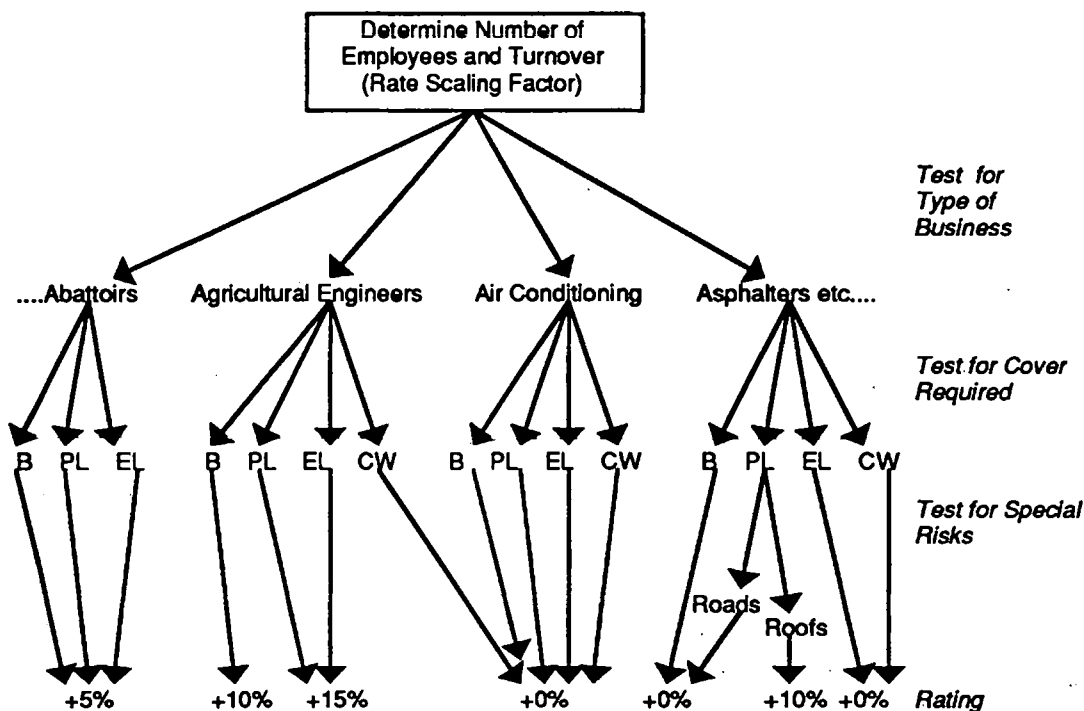
Tests at intermediate nodes prune the search at lower nodes. For instance in the small sample of commercial undertakings shown above there are no "Contract Works" (CW) risks to be considered for Abattoirs. This is because unlike Air Conditioning companies, Abattoirs do not undertake off-site work on behalf of other companies. Thus, for Abattoirs, the search tree is pruned so that such risks are not considered. Similarly, the tree can be extended as a result of tests at intermediate nodes. Thus asphalters are subject to a test for special risks (ie working at heights on roofs) in respect of Employee Liability (EL) risks. They are also subject to Public Liability tests (ie endangering the public, when they are working on roofs (dropping tools) and on roads (causing vehicle collisions).

The underwriter used the freehand quotation "prop" to gradually develop the quotation. As he worked through the construction of the quotation, from type of undertaking, then to cover required and then to actual work proposed, the pruning of the search space (ie limiting the next set of factors to be considered) plus the incorporation of special tests was very evident.

Diagrammatic Representation of Problem Solving Components

The diagram is the first and most simple of those that we will use to represent the components of problem solving of this and the following tasks. The headings at the top are standard to all diagrams and the arrows show which of the major components are present. In this case we show that a single domain is searched leading to a straightforward primary conclusion. To some extent, because of its simplicity, this diagram is rather redundant. However, it is included because of its relevance as the most simple case in comparison with some of the more complex architectures which follow.

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Single Domain

=> => => => => => => => => => => ▓▓▓▓▓▓ => => ▓
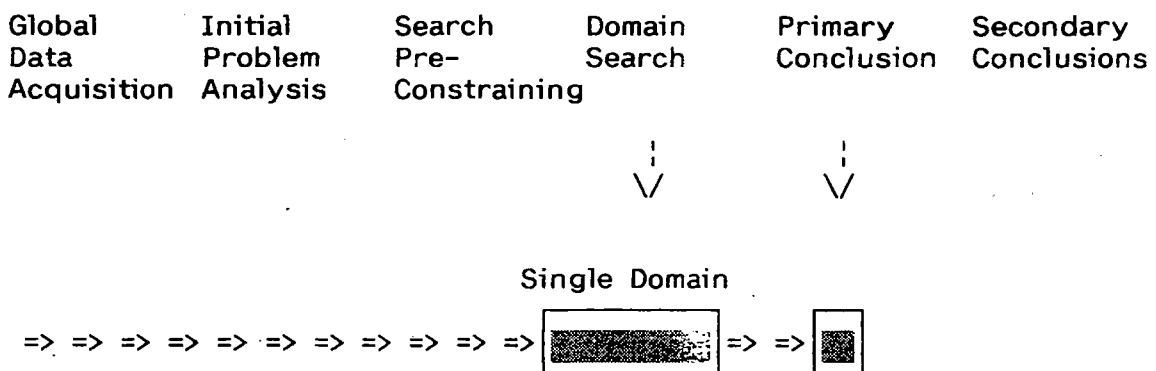
Commercial Implementation Details

Following performance problems with IBM's mainframe ESE shell, described more fully in Chapter 6, the system was successfully implemented on an IBM 3090 under IMS using Aion Corporation's ADS. The author had only a small involvement with the eventual implementation of this system. However, a key requirement from the outset was to support the "virtual form" (as described above) used by the underwriter and to access the risk tables which had been accumulated over a long period.

The "virtual form" records the key elements of the risk and highlights interactions leading to special risks. For instance, the fire risk to a timber merchant (business activity) is considerably increased if the business is located next to a fish and chip shop (business location). Similarly, the risk rating for a building site will be considerably increased if a crane is working on the site since the risk tables show that cranes (business equipment) often drop items on the heads of the site operatives (business staffing).

The system had to prepare a text based justification of the premium using the same style as the virtual form used by the underwriter. Without such a justification, which was always prepared by the human underwriter, the broker was unable to justify the premium to the customer and get his business.

The expert underwriter used for this system was very cooperative but was convinced that his expertise could not be replicated. Until it was pointed out to him he did not realise that he was approaching every risk on a very similar basis and, although each handwritten "form" was constructed from

scratch, he was actually working through a well-established and reproducible process.

Additional Comments

This task involved heavy use of tabular material and "rule books" and was a good example of where the success of the final implementation depended on encoding and integration of tabular and rule data.

3.3   FOREIGN EXCHANGE EXPOSURE

Sponsoring Organisation

The investigation and commercial implementation were sponsored by the Corporate Treasury Department of a major aero engine manufacturing company.

Introduction

The manufacture of major "bespoke" products, such as aero engines, requires the purchase of many raw materials and sub-components over a period which may extend for several years from the signing of the manufacturer's contract with the purchaser. Many of these sub-components and raw materials will come from outside the UK and the manufacturer faces considerable risk from changes in the sterling exchange rate against the currency of the supplying country.

Corporate Treasury departments have a duty to advise purchasing departments in operational divisions how they might best hedge against such currency changes. If this task is not undertaken the profit from the contract, often negotiated against tough competition, can be easily eroded during the years between contract signature and product delivery.

This task originated from a company wishing to extend the expertise of its Corporate Treasury staff to purchasing units so that such units could deal with the mass of small enquiries on this topic without referral to the Corporate Treasury Department.

Author's Role

The author instigated this expert systems project, undertook the primary knowledge acquisition and collaborated in the core design of the implemented successful expert system. He did not undertake the detailed design or implementation.

Outline Task Description

The expert in the Corporate Treasury Department is usually called by telephone to advise on a strategy for containing foreign exchange exposure on a particular purchase that is in course of negotiation or finalisation. Advice is sometimes given by telephone or, more usually, by facsimile. Such advice is given as a form of wording to be placed in the purchase contract or, if a contractual solution cannot be negotiated, the foreign exchange hedging strategy to be adopted by the purchasing department.

Form of Investigation

Knowledge acquisition was conducted with two key experts. Although the experts were held in very high esteem by their company, their knowledge was not particularly complex· or extensive and all relevant knowledge was acquired very quickly. As proposed below the expertise can be represented by a simple search tree and implementation was exceptionally straightforward.

Findings on Nature of Problem Solving

A simple search is conducted. The expert first ascertains the anticipated exposure in size and currency. If the amount is small or if US dollars are involved then no action is required. In other cases the enquirer is asked to attempt a number of strategies that place the risk on the supplier. For instance if the supplier will bill in sterling or US dollars then no future exposure risk is involved.

If the supplier will not adopt sterling or US dollars then acceptability of another major currency (eg German Marks) is tried. If the supplier will bill in this currency then hedging against the future risk is routine and quite straightforward. If the supplier insists on billing in an obscure foreign currency (eg Turkish Lire) then the purchasing unit is provided with a bespoke hedging package arranged by the Corporate Treasury through the company's bank.

Problem Structuring Props (Tools and Aids)

A standard pre-printed A4 form is used to record the purchase exposure details and fully supports the expert during the progressing of the enquiry to its completion. The form is an ideal example of a problem solving prop and is simulated in the expert system input screen and report.

Generic Data

Reuters and Telerate futures pages are used extensively to determine likely future exposures and, by future purchase options on foreign currency, to

set in place the hedging strategies where necessary.

Specific Data

Requests from purchasing units are usually followed by a short series of questions to determine the specific problem data. This is usually very simple and consists of the amount of the purchase, the currency involved and the likely contract and payment dates.

Knowledge Sources

A single knowledge source supports the straightforward problem solving algorithm.

Open/Closed Problem

The problem is closed with no opportunity for novel problem descriptions or solutions. If exposure to a very obscure currency is anticipated then the "expert" merely devolves the hedging strategy to the company's bank. Thus, as with Pensions Management, any problems outside the closed set are referred outside the competence of the expert and, in the implementation, outside the expert system.

Proactive/Reactive Problem

The problem solving is solely reactive to requests from purchasing departments.

Problem Solving Architecture

Jointly with Commercial Business Underwriting this task represents a very simple form of search of a single knowledge sub-domain. The search is more-or-less algorithmic and can be represented by the simplest expert system shell. Straightforward backward chaining search was chosen for this implementation with discrimination tests performed both at intermediate and terminal nodes.

However, as with Commercial Underwriting, the size of the search space was very large. Whereas the Commercial Business Underwriting the search space was very wide but shallow, the search space for Foreign Exchange Exposure is both wide (as a result of number of currencies for which exposure must be protected) and also, in the extreme, deep (as a result of complex hedging strategies occasionally used).

The system is used partially to advise telephone callers and not all the information is available in advance. The user is thus provided with prompts throughout the search to ask the caller for his/her preferences on currencies and on the likely flexibilities of supplier to accept payment terms in dollars or German Marks.

Diagrammatic Representation of Problem Solving Components

We show Foreign Exchange Exposure as having the same simple architecture and components as Commercial Business Underwriting.

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Single Domain

=> => => => => => => => => => => | ▓▓▓▓▓ | => => | ▓ |

Commercial Implementation Details

This simple system was implemented using Crystal 3 where elaborations in the user interface compensated for the straightforwardness of the knowledge. The system was required to interface with an invoicing package but this was relatively easy to undertake and no problems were encountered.

Additional Comments

This investigation and subsequent implementation proved most useful in demonstrating the power of simple expert systems to support the expertise of quite senior executives. However, the investigation is recalled with some embarrassment since it was initially felt by the sponsoring company that the

individuals concerned were the custodians of very special expertise.

The author's failure to find anything except simple algorithmic search caused some difficulties with the sponsor and, in fact, precluded the investigation of further tasks in case they too were found to be mundane.

52

3.4   BENCHTOP ELECTRONIC DIAGNOSIS

[The author's account of the commercial management of the implementation following this task investigation, but not the detailed research findings, is included as a chapter in Bramer (1990).]

Sponsoring Organisation

The commercial aspects of this investigation were undertaken with the permission of British Telecom who later sponsored the resultant full-time system developments.

Introduction

This task differs from Telecom Repair Service (reported later in this chapter) in that diagnosis is conducted on equipment that can be fully accessed for testing purposes. In Telecom Repair Service the diagnostic domain was a remote local line network where hypotheses could only be positively tested at some considerable manpower expense. In this task the test domain is a piece of electronic equipment isolated from its normal environment and available on a technician's bench for full diagnostic testing and repair.

Author's Role

The author was solely responsible for instigating this investigation,

undertaking all knowledge acquisition and analyzing the acquired knowledge. He was jointly responsible for the design specification of the TRACKER expert system shell but not for its Prolog implementation.

Outline Task Description

Faulty electronic equipment is removed from telephone exchanges, either public or at customer's premises, and brought to a local workshop for repair. The repairs are undertaken on a diagnostic technician's bench with the aid of instruments, including meters, oscilloscopes etc, and on completion the repaired units are returned to stock for subsequent field use.

The technician's undertaking the task each have a range of individual pieces of equipment with which they are familiar. In the case of Benchtop Electronic Diagnosis the individual technician was expert at repair of a switched-mode power unit which supplied the various DC power supplies to a medium-sized customer switching system ("PABX").

Form of Investigation

The investigation of this task chronologically followed Telecom Repair Service reported later in this chapter. In the Telecom Repair Service investigation knowledge acquisition by concurrent verbalisation had been used. However, Benchtop Electronic Diagnosis was not amenable to concurrent verbalisation due to the extended time undertaken for diagnosis. Much of the technician's time is actually involved in disassembly, setting up test conditions, component replacement and final assembly and testing. In some

54

cases a technician can be involved for a complete day on the diagnosis and repair of a single unit. Much of the time involved has little problem solving significance and concurrent verbalisation would be inefficient in the extreme.

As a result, this task provided the author with his first major opportunity to undertake knowledge acquisition using the types of knowledge acquisition techniques discussed more fully in Chapter 4.

The knowledge was initially analyzed and a rule set developed on a "paper basis". This course was adopted because there was no convenient shell in which to implement the expertise and the author collaborated with his working colleague, Dr George Pollard, to design a purpose built shell, the TRACKER shell, which Dr Pollard implemented in micro-Prolog. In practice this collaboration required that the form of representation and the needs of the inference system were iteratively specified by the author in parallel with the ongoing shell implementation.

Findings on Nature of Problem Solving

Diagnostic search for this task has a heavy procedural element. In the default case, and after disassembly of the equipment, a series of tests are performed in a set order. The simple objective of such tests is to enable the diagnostician to encounter an abnormal test condition or diagnostic clue at the earliest possible opportunity.

Once such a disordered condition is discovered, the initial diagnostic procedure is suspended and a new procedure instituted on the basis of the discovered clue. This "clue-driven procedural diagnosis" continues

recursively until the faulty component is isolated.

The default method of testing described may be entered at different points depending on the nature of the reported fault in the equipment or obvious initial visual signs of broken or burnt components.

The diagnostic search may be very adequately described as proceduralised hypothesis-and-test with dynamic search tree rearrangement and pruning on the basis of discovered diagnostic clues. The TRACKER shell was designed to support this specialised, but quite straightforward, method of problem solving.

In particular the ability to encode procedural search without relying on contrived procedural reasoning through rules was particularly advantageous in the Benchtop Electronic Repair implementation and the several further commercial implementations using the TRACKER shell.

Problem Structuring Props (Tools and Aids)

The diagnostic technician did not use any form of problem solving prop for straightforward diagnoses. However, for more difficult or novel faults, he always made brief notes which summarised his current progress and the various tests which he had undertaken. In a similar fashion to the notes used in preparing quotations in Commercial Business Underwriting, these notes were structured in an essentially fixed format.

When confronted with a novel or ambiguous test result, the diagnostician would often sit away from his test bench and attempt to solve the problem

56

on the basis of the information recorded in his notes rather than on the basis of the overt meter and other test displays available on the bench.

The notes were recorded in a pocket book which the diagnostician referred to as his "diary". The presence of all information on the fault on a single page seemed to act as an "external memory". Consultation of the notes seemed on several occasions to be the essential trigger which enabled the diagnostician to progress the diagnosis to its successful conclusion.

When the fault was eventually isolated the notes became a source of reference for further diagnoses and were frequently consulted by the technician at times of difficulty.

Generic Data

The notes or "diary" as just discussed formed an increasing asset of generic data for the diagnostician. More formal generic data was provided by circuit diagrams, circuit descriptions and equipment manuals. However, as described below under "Commercial Implementation", many of these latter "official" sources of information were out-of-date or of little use for diagnosis.

Many of the faults encountered existed in new modified versions of equipment. Such modifications were rarely shown on the information provided for the diagnostician and thus his personal notes became an increasing source of valuable generic information.

Specific Data

Specific problem data was available from a Fault Report Docket completed by the field technician who had removed (and replaced) the faulty unit at the customer's premises. In many cases the dockets were vague and often did not mention the abortive local repairs which had been attempted and were the actual cause of more damage than the original fault. (For instance, the local technician might replace the fuse of a unit damaged by overloading and thus remove a useful diagnostic symptom for the workshop diagnostician.) In consequence such docket information was largely ignored by the diagnostician who would usually undertake a "default diagnosis" as if no prior information had existed.

Knowledge Sources

The domain can be considered to be served by a single knowledge source.

Open/Closed Problem

The problem is closed and restricted to a finite set of faults. With some few exceptions all faults can be mapped onto the failure of a single component. Therefore the number of feasible diagnoses is approximately equal to the number of components in the equipment.

Proactive/Reactive Problem

Diagnosis is undertaken purely on a reactive basis when the diagnostician is presented with a faulty piece of equipment.

Problem Solving Architecture

This task can be considered as a dynamic search tree with tests at both nodes and leaves.

In this form of search the tests at intermediate nodes are used to generate hypotheses which may be tested at subordinate local nodes in the tree. However it is also common for the results of tests at intermediate nodes to generate hypotheses related to distant nodes in the tree. Such unexpected "clues" may result in a search which "jumps around" the search tree. This form of search can be considered as a dynamically re-arrangeable search tree and, in implementation terms, can be considered as procedural search subject to forward chaining opportunistic control.

Shown on the next page is a simple example where a test at (*1) switches the search to (*2). It is to be noted that the tree below (*1) is not pruned and exhaustion of the hypothesis at (*2) will in due course allow the search to revert to the tree below (*1).

BENCHTOP ELECTRONIC DIAGNOSIS

Intermediate
Test at Node

(*1)

(*2) Search switches
to Distant Node

Final   Confirmatory   (Fault   Identification)   Tests   at   Leaves
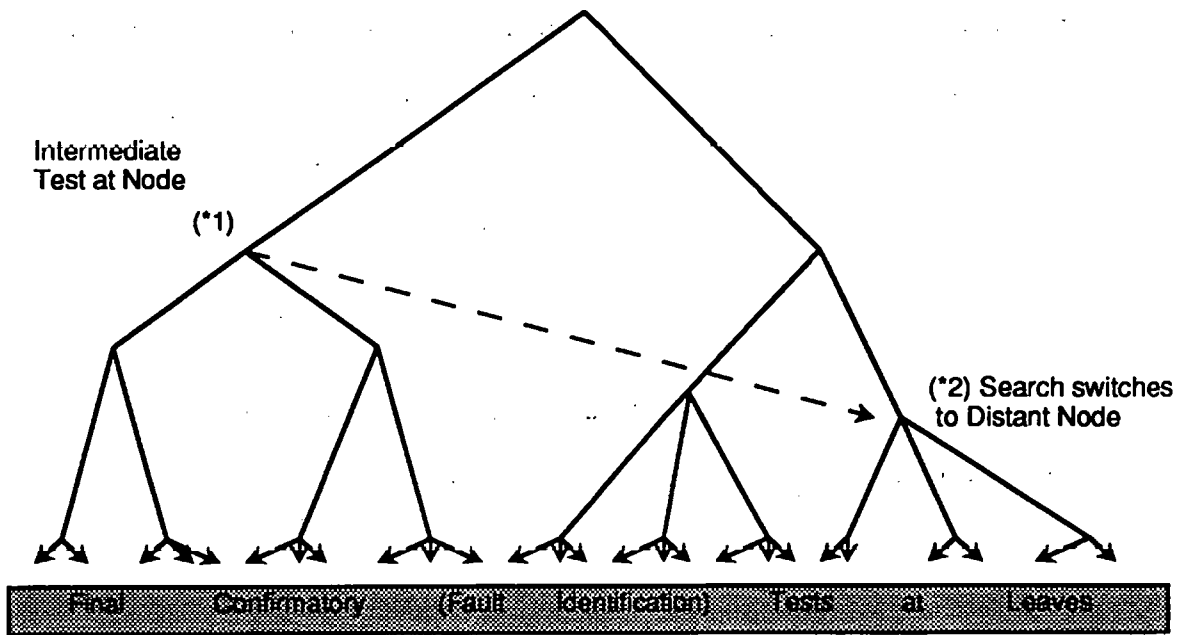
Diagrammatic Representation of Problem Solving Components

Again, although the actual search is more complicated than Commercial Business Underwriting or Foreign Exchange Exposure, the problem can again be shown diagrammatically as search of a single domain leading to a primary conclusion as to which component is faulty.

Global	Initial	Search	Domain	Primary	Secondary
Data	Problem	Pre-	Search	Conclusion	Conclusions
Acquisition	Analysis	Constraining			

Single Domain

=> => => => => => => => => => => => =>
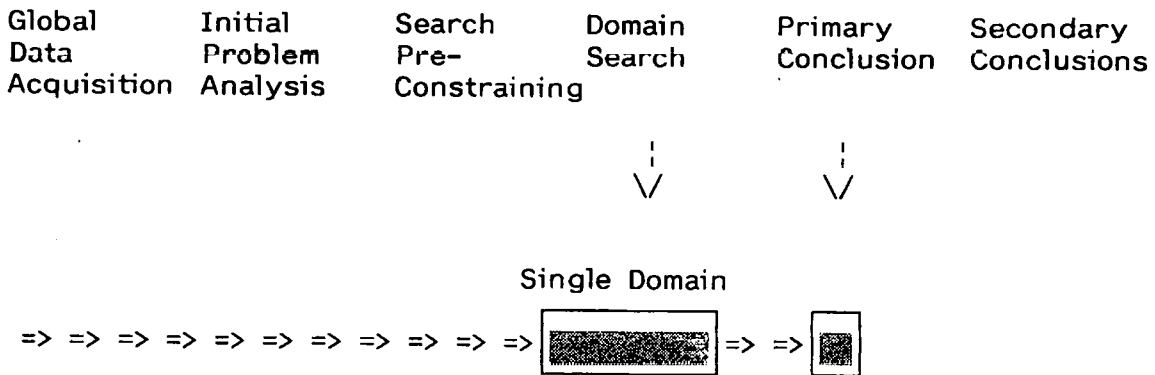
Commercial Implementation Details

The TRACKER shell was implemented by Dr Pollard and used for four internal British Telecom applications. It was later re-implemented by other colleagues, again with the author's design collaboration, and used by British Telecom for a further series of diagnostic expert systems implementations.

The TRACKER system was very successfully used by enthusiastic technicians at a number of BT sites and was the subject of an OU/BBC Alvey video case study. Its universal usage at a wide number of BT sites was prevented due to limitations in the hardware budget necessary to provide IBM-compatible PCs at every benchtop work position. (It is not practical to provide a benchtop diagnostic system unless it can be delivered at the actual working position which is the benchtop. Use of a shared machine at a repair centre requiring the technician to get up from his working position at every stage of the diagnostic testing results in the system not being used since it is easier to puzzle out the solution to the problem rather than conduct an inconvenient "walking dialogue" with the expert system.)

Although the hardware problem could have been overcome there still remain three other problems that would today prevent the wider utilisation of the system as then implemented. These problems are described immediately below.

Problem 1 - Cost of Implementation for Multiple Equipments

First, BT workshop technicians repair a wide range of "electronic boxes". The TRACKER system was initially used to diagnose faults on a switched

mode power supply and a pulse code modulation unit and, later, and with only my partial involvement, certain other "boxes" such as modems etc. Each of the expert systems took several weeks to implement and test even though, at this stage, there was considerable experience in the use of the TRACKER shell and the author and his colleagues were becoming very skilled at knowledge acquisition and expert system implementation. Nevertheless, the cost of this implementation effort was too great if fully funded by the repair service profit centre and multiplied by the wide range (several tens) of types of units repaired in each centre.

Problem 2 - Changing Equipment Specifications

Second, and a recurrent problem for many proposed engineering repair expert systems, was that the specification of equipment being repaired was undergoing continuous updating. Thus items were coming into repair with various circuitry changes and additional components, often to cope with some of the faults which had plagued earlier variants of the equipment. For instance, in the switched mode power supply system, a major component of the diagnosis was to detect which of fourteen identical diodes had failed in service. Use of the expert system (or the presence of the very skilled expert) prevented the effort of exhaustively unsoldering and replacing each of the diodes, sometimes leading to secondary damage, to find the defective component. In the later versions of the equipment all of the diodes had been replaced by a more robust component with the result that the focus of probabilities of the cause of the faults had moved to other component groups. Electronic diagnosis is often like looking for a weak link in a chain. If a particular link is strengthened then the probability of the cause of faults moves to the next weakest link. If a diode has been uprated then

a transformer becomes suspect in the uprated equipment; by uprating a transformer then a circuit track becomes the likely suspect etc.

The result of these ongoing circuit improvements, some of which were triggered as design changes resulting from the recurrent faults highlighted by the expert system, had a profound effect on the usefulness of the diagnostic expert system used for diagnosing the later variants of the "boxes". There was an urgent requirement to produce updated versions of the knowledge base to cope with the changes in the fault patterns of the variants. Heuristics in the expert systems match the fault pattern in that they lead the user in the most efficient route to discover the fault. If the fault probability changes then the heuristics will usually discover the new faults but in a very inefficient and roundabout manner. This means that the expert system is no good and will not be used more than once or twice since it will lead the user in a path that probably takes longer than unaided exhaustive search.

The only practical solution for this second problem is to provide expert systems that can be user updated when the subject of diagnosis changes its characteristics. This brings forward a wide range of problems ranging from the "learning" issue right through to user interfaces allowing users to make small changes to the knowledge base. All of these problems are substantial and very much outweigh the issue of providing a relatively simple expert system on inexpensive user hardware. This matter is further discussed under "Design-Diagnostic Cooperative Expert Systems" at the end of this section.

Problem 3 - Cost-Justification of an Expert "Teaching" System

The third problem was not actually realised by the system sponsors although the all those involved in TRACKER system implementation secretly realised it very early in the work. Simply put, many expert systems will not be used once they have "trained" the user to be an expert. In the first TRACKER implementation (which was the subject of the Alvey video) the author provided the initial trial system for a young technician who, very rapidly, was able to undertake the same level of repairs as the original expert. After exploring most of the diagnoses he did not use the expert system any more since he had effectively been trained by it. As a result, although the user was very pleased with the expert system and set great store by it, the system sat on the back of the bench for most of the time and was only switched on when the user encountered a tough problem for which the system had not "trained" him. At this stage the system was switched on, a successful diagnosis undertaken and then not used again until the next unusual fault is encountered. This is exactly the same as the human "master and apprentice" model where the increasingly experience apprentice asks questions of the master with decreasing frequency until he is fully trained.

Of course, the use of expert systems as training aids is very worthy. However, it does not impress management who expect to provide a resource only if it is cost-justifiable on the basis of its long-term continuous usage or as an aid to overall increases in worker productivity. An expert system that, once it has trained its user, becomes largely unused, does not fit this model. Although the TRACKER sponsors never fully realised this problem, it would have undoubtedly surfaced later and the traditional trainers, with a vested interest in maintaining expensive training courses, would have

pounced upon this issue and criticised the expenditure on infrequently used hardware and software.

Further, it would have been impossible to justify the expert system on the grounds of a productivity increase. There were no apparent increases in productivity because, if anything, usage of an expert system slows down a worker. Once self-experienced, or trained, there seems to be no more productivity benefits that can accrue to an expert by use of an expert system to support a straightforward task such as Benchtop Electronic Diagnosis. In the cases of some tasks discussed later, such as Multinational Tax Planning, use of the resultant expert system could aid the productive solution of complex problems. However, such justification was not possible for Benchtop Electronic Diagnosis or any of the other unreported similar electronic diagnostic tasks investigated commercially by the author.

The form of diagnosis found in Benchtop Electronic Diagnosis is expensive and difficult to support. Nevertheless, the knowledge may be acquired reasonably easily and additional work needs to be done to overcome some of the above practical problems so that diagnostic expert systems, perhaps coupled with ATEs (Automatic Test Equipments), may become more commonly used.

Design-Diagnostic Cooperative Expert Systems

A particular barrier to commercial implementation, that of changing equipment specifications is discussed above. This issue is now further discussed in the context of the lessons learned from this task in respect of inter-expert cooperative problem solving.

In effect there is a feedback loop between repair diagnosticians in the field and headquarters designers of electronic equipment. (This applies equally well to domestic TV sets as for the telecommunications equipment that was investigated.) The first version of any new piece of equipment is solely the product of the designer's work but later versions and modifications depend on experience of the equipment in use.

All electronic equipment fails and statistics are gathered on components to determine those that are the key causes of failure. In switched mode power supplies the main diodes are often underspecified, for reasons of economy, and have a high failure rate. Later, or modified, versions of the equipment have higher specification diodes and this source of failure becomes far less frequent. As already described this has the result of changing the heuristic search pattern. Circuit diagrams for the modified versions of the equipment are produced and these diagrams, together with component listings, are available to the repair diagnosticians when the modified equipment comes in for repair.

In theory this feedback process is relatively sensible. However, there are major problem due to the fact that the "design expert" (or small team of designers) and "repair expert" (of whom there may be several hundred at different sites) do not communicate in any effective sense.

First, the designer does not really know what causes the equipment to fail. He receives data on component failures but this is only partially useful. For instance, the same capacitor may be used twenty times over in the same equipment but it is only at one location in the circuit that it is overstressed and fails. And the failure may not be due to an inherent fault in the capacitor but due to the fact that field installation technicians regularly

make accidental short-circuits of the equipment output during the installation testing. What is really required is a circuit modification to give short-circuit protection to the equipment rather than simple replacement of the capacitor.

The diagnostic technician knows the causes for circuit failure but there is no mechanism for him to convey these back to the designer. Therefore this half of the feedback loop is grossly defective but could be radically improved if communicating expert systems were used by both the diagnosticians and the designer. In this case the designer would benefit from knowledge of the actual causes of faults rather than the raw component statistics.

The feedback from designer to diagnostician is also a serious problem. Diagnosticians encounter many circuits, probably the vast majority, where they are only provided with a circuit diagram and component listings. In later versions of circuits there are many modifications with single or small groups of components added on to the main circuit board. Usually the diagnostician has only a vague idea of the purpose of these extra components and, quite often, he does not know what subsidiary circuitry in the main equipment is designed to achieve. For instance, in the pulse generation circuitry of power units, the diagnostician was unable to describe the purpose of many of the components. However, he knew that they all had a purpose and could describe those most likely to fail. The situation is similar to medicine where a surgeon knows that a problem in a particular gland is the source of an illness but has no idea of the physiological connection between the defective functioning of the gland and the symptoms.

This situation could be radically improved if the diagnostician had access

to an expert system including the design knowledge and circuit objectives of the designer. On a more simple level, much could be done by obliging the circuit designer to produce circuit design notes in advance of the introduction of the new equipment into the field. However, in many organisations, the production of such circuit notes is the responsibility of a Training Department and their availability lags well behind the appearance of the equipment in the field. By the time that they appear the repair diagnosticians have already painfully discovered (or confused themselves) as to the actual way in which the circuit operates.

These and similar issues are further considered in Section 6.10 under the heading of "Opportunities for Communicative Expert Systems".

3.5   AUDIT WORK PLANNING

Sponsoring Organisation

The investigation and extensive work on the implementation of the resultant expert system was sponsored by Coopers & Lybrand's audit practice.

Introduction

This very successful expert system was implemented by the author's Knowledge Engineering Group at Coopers & Lybrand.   The implementation was difficult and protracted and in a professional area not previously addressed by expert systems technology.

Author's Role

The author's research role was in the core knowledge acquisition, audit task analysis and system design stages.   Although he had oversight management and responsibility for the later stages of the implementation, much original and very ingenious implementation was done solely by others, particularly to enable a very large and comprehensive expert system to be run within the then limitations of a portable PC.

Outline Task Description

Auditor's are primarily responsible for providing an opinion as to the

validity of the statutory accounts of companies. Unlike fraud investigators, who make a complete check of all financial documentation, auditors make an examination of a statistical or representative sample of the firm's financial records. These may be either paper records, such as invoices or stock records etc, or financial and other data held on the firm's computer.

The task of the audit manager, whose expertise was targeted for this system, is to plan the workload of the actual staff who make the physical checks of financial records at the company's premises. Audit staff time is an expensive resource to accountancy firms and the audit manager's overall goal is plan a professionally competent audit entailing expenditure of the minimum staff effort.

The planning is traditionally completed manually and it was known that managers often adopted a 'belt and braces' approach so that the resultant audit was 'safe' but over-resourced. The actual output from both the existing manual process and the expert system is a multi-page workplan of tasks for audit staff to undertake.

Form of Investigation

Accountants have a long-established tradition of producing competent manuals on audit practice. These exist both to give material for the tough accountancy examinations and also, in theory, to guide the day-to-day work of audit staff. In practice the actual expertise of conducting an audit, and the particular skill of audit work planning, is learned "on the job" and maintained largely in the heads of the practitioners. Moreover, from our initial investigations, it seemed that there was considerable difference

between the knowledge and views of different audit work planning 'experts'.

As a result, and as an aid to efficiency, a particular audit manager was selected as primary expert and his knowledge used to design the core system. Later, it was decided to use this expert, by then familiar with the expert system process, to seek out detailed knowledge from his colleagues so that this could, where appropriate, be incorporated into the system. Nevertheless, the initially discovered form of working, as reported here, was merely fleshed out in detail by the use of the expert as an intermediary to the wider knowledge of his colleagues.

Initial knowledge acquisition was carried out by working through the files of established audit plans with the expert using the particular cases to explain and justify the resultant work plan in the file. One of the primary side-effects of this process was that, unless the expert had actually prepared the plan, and could remember the details, it was often very difficult for him to provide a retrospective justification for the decisions made in the file.

In many cases particular 'audit steps' (ie checking processes within an audit) were inexplicably not carried out, carried out to no obvious purpose or carried out on material that had already been checked in another part of the audit.

This investigative work, carried out at a comparatively early stage of the project investigation, provided considerable further managerial justification for the need for an expert system to 'discipline' the audit work planning process.

Knowledge acquisition sessions were audio recorded. Although in a 'balance sheet domain', where video recording would usually be useful, the planning process is at a higher level than the actual scrutiny of figures and it was possible to acquire knowledge without extensive reference to tables of figures etc. Nevertheless, copies of the papers from audit files were used in the analysis of the knowledge acquisition tapes.

Findings on Nature of Problem Solving

Knowledge acquisition revealed that the audit process depends on a very considerable amount of detailed knowledge. Nevertheless the task is highly segmented and can be considered as nearly algorithmic in nature. The audit manager first collects basic details on the particular client (eg size of operation, locations, computer systems etc) and also consults the audit file for the audit partner's overall strategy for the audit.

Using these basic facts the manager considers each of the audit areas (eg Creditors, Debtors, Stocks, etc) in turn so as to select and decide the level of severity of tests which best fit the overall strategy and the precise physical characteristics of the firm and its records. At the end of this process the manager considers whether any combination of factors from different areas (eg abnormally high stocks, high debts but apparently normal turnover) make it necessary to devise any special tests or strengthen (or combine) any tests from particular audit areas.

Problem Structuring Props (Tools and Aids)

The audit manager uses a pre-printed form which includes a schedule of

possible audit tests etc. The form is completed by the manager who is responsible for filling in details of tests to be performed together with their level of severity. For instance a very severe test may involve examining every record and a minimally severe test may just involve a few samples. The form is often heavily annotated with details of special, abbreviated or combined tests. Explanations for the annotations are usually entered on the form in an abbreviated cryptic form. The form plays a very significant role in structuring the audit manager's planning task and, in the implemented expert system, is reproduced as part of the system output.

Generic Data

The company audit manual is the sole source of information used by the audit manager. It is only very minimally helpful in the task and, except for background information, can be sensibly ignored.

Specific Data

The 'Audit File' of the company whose audit is planned is used as a specific input to the audit planning process. It contains details of previous years' audit schedules, which are sometimes used as a guide to the current year's plan, together with background details on the company. Sometimes also included in the file are details of pre-audit discussions which indicate any major changes in the company's trading pattern.

Knowledge Sources

Separate knowledge sources can be envisaged for each of the separate audit areas. Because of the low level of interactivity between areas the knowledge sources can be applied singly in a linear progression matched to the production of a report. The interaction knowledge source (implemented in the expert system as a "conciliator" module, the term "conciliator" coming from use by auditors of the word "conciliation" to mean resolution of final problems between auditor and client) is applied last and looks for particular 'dangerous' combinations of facts and conclusions from the individual audit areas. This section is implemented as a miscellaneous collection of forward reasoning rules sensitive to particular combinations of conditions that may be evident following completion of the main part of the inferencing. Such rules detect, for instance, combinations of cashflow and work-in-progress conditions that could lead the auditor to be suspicious of an otherwise apparently solvent firm.

Open/Closed Problem

The problem is closed since there are only a finite number of audit tests and combinations of tests that may be planned.

Proactive/Reactive Problem

The task is solely reactive to the annual statutory need to produce an audit. The audit manager is never prompted to undertake an audit by, for instance, unusual combinations of financial symptoms from the company.

Problem Solving Architecture

This is essentially a series of independent search trees each related to a sensibly separate part of the domain. For instance, when planning the part of an audit covering a firm's Stocks, the audit manager would make a plan for his staff to examine completely different areas of a firm's records and physical situation from those which he would plan to examine in an audit of Debtors or Trade Creditors.

Nevertheless, although he is planning an essentially separate series of audit examinations, each of the different searches actually relates to the same underlying company. For this reason such multiple independent searches may be preceded by the collection of "global value data" which provides input of common variables (eg firm's trading activity and size, form of computer records etc) to each of the multiple searches. This information is used as a common pool for the plans of each different part of the audit. For instance, the fact that a firm is a wholesale jewel merchant will be a key input to how Stocks should be assessed and what is normal for Debtors in the jewellery trade.

Similarly, there is a very limited requirement to make overarching inferences based on the findings of more than one of the multiple searches. For instance, a conclusion during the planning stage of the full audit that there was an unusually large Trade Credit situation coupled with excessive Stocks would point to a firm that should undergo further scrutiny in respect of poor results (ie large stocks and debts to suppliers means that expected trading had not materialised). Such overarching inferences are usually made by the Audit Partner (ie the person who signs the statutory audit) and will affect the nature of the plan devised. However, in planning the

75

audit, the audit manager must anticipate where and when such conclusions will be made. In making his plan the audit manager must provisionally examine the financial situation of the company being examined and thus the architecture of audit planning is more-or-less identical to audit itself.

The final overarching conclusions are represented by a "conciliation" section which, in the final Audit Work Planning expert system was implemented as a "conciliator" part of the inferencing system.

It is to be noted that audit planning is an iterative process with the manager often being unsatisfied with the "final plan" for relatively trivial reasons. He will then feed back his findings as data for a new plan which is more likely to be acceptable to himself and his Audit Partner.

Audit Work Planning as described above can be illustrated as on the opposite page.

AUDIT WORK PLANNING

```
                         GLOBAL VALUE DATA COLLECTION

   STOCKS        DEBTORS        TRADE CREDITORS        ETC        ETC

        P   R   I   M   A   R   Y       C   O   N   C   L   U   S   I   O   N   S

Iterative
Feedback          OVERARCHING CONCLUSIONS - "CONCILIATION"

                         SECONDARY CONCLUSIONS
```
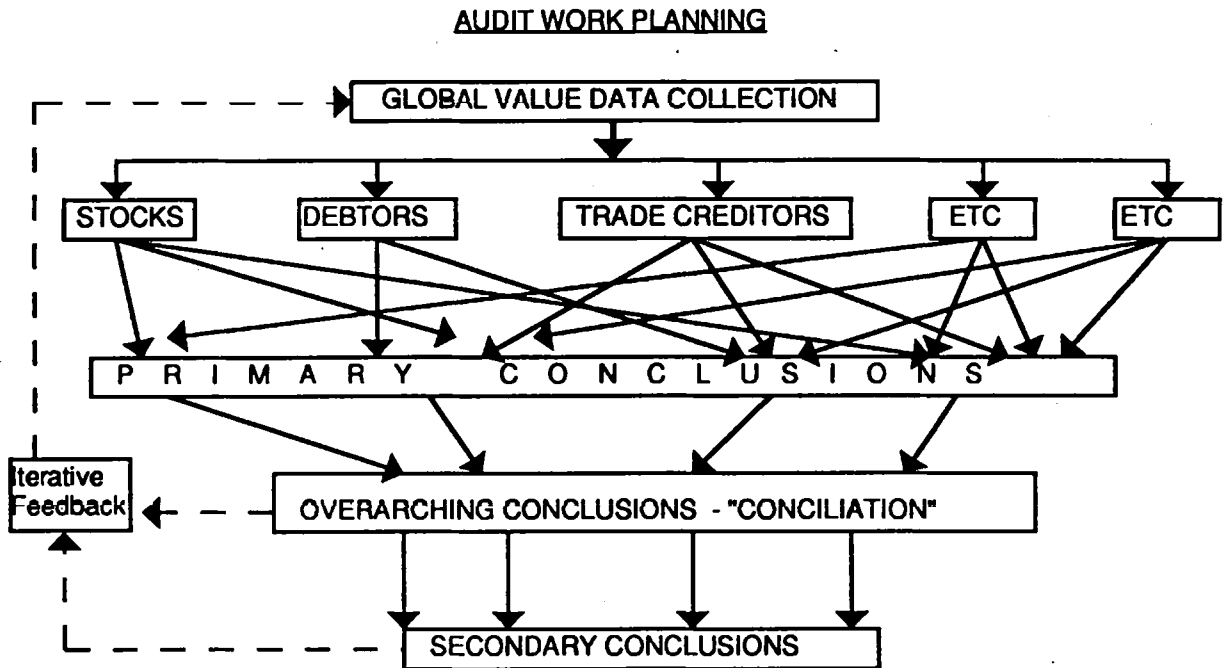
Diagrammatic Representation of Problem Solving Components

The diagram now shows a significant change to the simple single domain search illustrated for the previous cases. A global data acquisition component is shown feeding initial data into a number of separate sub-domain searches. (In this and the following cases just three or four multiple sub-domains are shown although in practice a greater or lesser number of sub-domains may be searched.)

In this example all sub-domain searches are independent, with no inter-sub-domain references (eg there is no requirement when planning a Stocks audit to refer to any initial or inferred information about Debtors), leading to

77

separate primary conclusions.   As explained above the outcomes of the
separated searches may lead to some "Limited Overarching Conclusions"
which are shown as a component on the diagram.   A feedback loop leads
from such overarching conclusions back to the initial problem data to
facilitate iterative problem solving.

Multiple Simple Search with Limited Overarching (Conciliator) Findings

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Multiple Domains with no Reference Linking

Limited Overarching Conclusions

Iterative Feedback
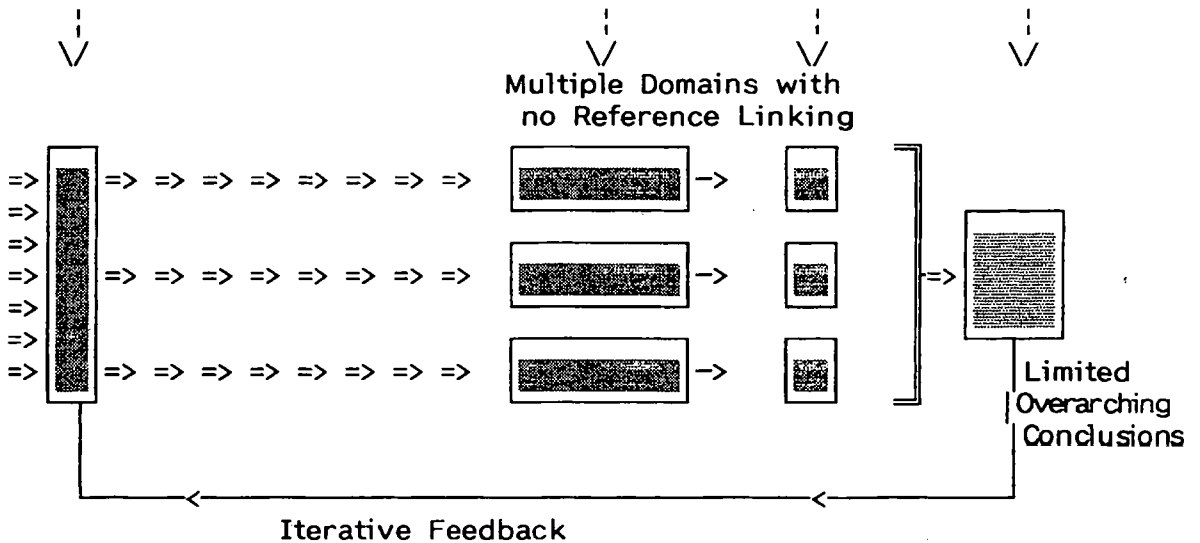
Commercial Implementation Details

The system was implemented as "Expertest" using a re-written version of Qshell, an in-house Lisp-based Coopers & Lybrand shell with particular facilities for providing fully-formatted reports. The system was delivered on a portable IBM-compatible PC for field usage. Very considerable ingenuity was required by the author's colleagues to deliver the system within the limited hardware environment required.

The system was commercially very successful and has remained in use over a long period. It is therefore useful to consider the following factors which were principally relevant to its success.

It was difficult and time-consuming to sell the system concept to the senior audit partner but, once the author had his "buy in", he was willing to

persuade all the other key audit staff that the system should be constructed and put fully into operation. This ensured that budget was available for the whole project and that the time of the key audit experts was made available.

System delivery issues were vital and were the principal key to the expert system's success and widescale usage. The system was to be used by audit managers working "in the field" at the sites of audit clients. All such staff carried portables PCs for routine spreadsheet and word processing usage. The system had to be delivered "hardware cost free" on these machines if it were to stand any chance of widescale adoption. (In the event some of the target machines required a memory upgrade to run the expert system but this incremental cost was insufficient to prevent the project going ahead.)

The system had to be very user-friendly and pose its questions to the user in a familiar form. Audit managers traditionally fill in a an on-site form which they take back to their audit partners (ie senior managers) to plan the audit. The expert system had to allow managers to answer similar questions to those on the existing form and to tick boxes in the traditional way.

The system output had to be a well printed and formatted audit plan. Traditionally audit partners would prepare a draft plan in nearly-illegible handwriting which a secretary would transform into an elegant document for discussion and agreement with the client. The expert system had to prepare exactly the same document which was "signed off" by both the client and the staff concerned with the audit. Cosmetic details of appearance (eg how it should be signed and the exact text in accordance with auditing

standards) were given great attention.

The system produced a paper based final audit report and, very strangely, the system was not fully accepted by all senior staff until it produced output on exactly the same yellow coloured paper as was used for the traditionally typed reports. Even more peculiar was the fact that, once the system had been widely accepted, the audit partners committee were so enthused with expert system planning of audits that they changed the colour of the paper to green. The result was that audit files then had a green streak if viewed from the edge (rather than the traditional yellow) and this was regarded as a sign of progress and to be admired.

In summary, after a very difficult "sell" supported by a friendly "inside champion" this system was accepted by the audit bureaucracy who eventually fully adopted it as their own and embellished it with the traditional trappings of their trade.

Additional Comments

The Expertest implementation and the effort that was involved in it demonstrated that competent expert systems often require more work on the peripheral delivery features, such as interfacing and performance enhancement, than the core efforts on knowledge acquisition, analysis and knowledge base implementation.

3.6   BUSINESS PERFORMANCE ASSESSMENT

Sponsoring Organisation

The commercial aspects of this investigation and the subsequent successful implementation were sponsored by the Department of Health.

Introduction

Government departments and large commercial organisations often suffer from the inability to respond quickly to business problems and opportunities occurring at their many geographically distributed operating sites. Further, in such large organisations, the quality of decision makers is usually high at the headquarters of the organisation (and perhaps at some regional centres) but there is usually less ability and experience at the organisation's periphery. Nevertheless, it is at the remote periphery that the organisation's eventual profit, loss or usefulness is actually decided.

In this investigation the Department of Health had an established "performance review" procedure where local decisions were annually reviewed by headquarters experts. A series of "performance indicator ratios" covering such aspects as staffing ratios had been established to enable the headquarters experts to undertake their task without problems arising from the absolute size of the underlying figures. On this basis it was believed possible to compare the efficiency of Health Areas which had different population sizes.

Further, very able staff in the Department's OR Department had attempted to construct a decision tree, based on evaluation of the performance indicator ratios, in an attempt to "automate" (on paper at least) the evaluation process of the headquarters experts. The requirement for an expert system was initiated when it was realised that the expertise involved could not be set down as a simple paper-based decision tree.

The expert system required, and finally supplied, enabled a rapid evaluation of the performance of any operational unit in comparison with the performance of all other similar units throughout England & Wales.

Author's Role

The author developed the concept and outline design for the expert system and undertook the initial feasibility studies and knowledge acquisition. Nevertheless, although managerially responsible, he relinquished day-to-day involvement in the major implementation effort which resulted in the final expert system.

Outline Task Description

The task, which in its original form took many hours, was undertaken by performance analysts and involved a thorough analysis of related groups of performance indicator ratios followed by the writing of a detailed report. The report covered a series of topics (eg personnel, buildings, services etc) all related to performance and then attempted a global conclusion where common factors were highlighted from the specialist topics.

Form of Investigation

The task was already well understood by the OR expert who had attempted the conventional investigation of the expert analysts' task. For reasons of expediency, and a wish to co-operate with and train the OR expert in expert systems techniques, the OR expert was used as a substitute for the actual analysts who prepared the report. This "vicarious knowledge acquisition" worked extremely well and it was not until a later revision of the implemented system that analyst experts were involved to any appreciable extent.

A pilot implementation was commenced before completion of knowledge acquisition for proof of concept and funding reasons.

Findings on Nature of Problem Solving

Although it had not been possible to express the problem solving as a paper-based decision tree, there is little doubt that the problem solving, as developed by the human experts, is almost entirely algorithmic. This may be because, as a government department, it was necessary to construct any reports on a logical basis for the purposes of public accountability and fairness.

Any "creativity" in the decision making seems to have been restricted to the choice of text for the report and, to some extent, to the points emphasised in the global conclusion section.

Problem Structuring Props (Tools and Aids)

The task was conducted on a very methodical and logical basis with considerable access to reports and standards for performance indicator ratios. The primary prop appears to be the well-established form of the final report which seems to have been the goal and organising structure to which all the bottom-up analysis was directed.

The flowcharts, devised primarily by the OR expert, appeared to be based on fragments of flowcharts used by the analysts themselves and these may themselves be considered as low level props to problem solving. It was interesting to note that the analysts' flowchart fragments were used as routine aids to construct the detailed findings for the report and that these, in turn and without the use of any explicit flowchart, were used to construct the overall conclusions.

Generic Data

Performance indicator relationship charts, devised by external professional advisers, were the primary source of generic data to support individual case analyses. Standard targets for performance indicator ratios, also externally set, were also employed. Again, as described above, these prescribed relationships and targets were used on an almost mechanical basis to construct the detailed findings of the reports.

Specific Data

For each analyzed operational unit, there was a set of performance indicator

data plus access to the previous year's data for comparison purposes. This data was all held on spreadsheet files.

Knowledge Sources

It is not realistic to consider the task knowledge as separate knowledge sources since, although the analysis covers different sub-domains of operational performance, the same analytical expertise is employed for all sub-domains.

The concluding analysis is similarly heavily integrated with the sub-domain analyses and the knowledge involved is essentially similar to that used for analysis of the sub-domains.

Open/Closed Problem

The problem is closed with no possibility of novel situations occurring. Input performance indicator ratios can vary widely but are reduced to figures in the range 0%–100% in percentile comparison with other operational units. Range values for three ranges (High, Medium and Low) within the 0%–100% band are set and no combination of inputs can be outside these ranges. Accordingly there is a large but finite set of possible combinations of ratios. Since ratios are only assessed in comparatively small groups there is no chance of massive combinatorial effects and all eventualities can be covered.

Proactive/Reactive Problem

The task is proactive to Departmental need to report. In practice it is reactive to the need to analyze incoming sets of data requested by the department.

Problem Solving Architecture

As in Audit Work Planning the experts for this type of task also undertake a series of more-or-less independent searches (best considered as "analyses") in different areas of the domain. For instance, in Business Performance Assessment the domain was naturally subdivided into a series of different health care disciplines (eg Cardiac, Obstetrics, Psychiatric etc) and the "business performance" in these sub-domains was separately assessed. However in Audit Work Planning the results of the sub-domain analyses, the primary conclusions, were the most important output of the task. The secondary conclusions, as a result of "conciliation" were of an opportunistic kind and of genuine secondary importance.

In Business Performance Assessment the reverse is true. The primary conclusions have only limited value in isolation and are primarily used as the "building block" evidence for the secondary conclusions. Such secondary conclusions, which deal with major "global topics" such as professional staff to manual staff working ratios, funding levels, customer service performance etc are based on a summation of the evidence found in the individual sub-domains. In fact there is considerable similarity between the findings in the individual sub-domains. They are only analyzed in isolation since their individual characteristics do not make it possible to sum

88

the complete "business" in one stage. The problem solving is thus described as "Multiple Simple Search with Secondary Grand Conclusions".

Thus "level of customer service" in a medical unit would be derived from statistics based on, for instance, length of waiting time for non-essential operations (eg varicose veins). In a geriatric sub-domain "level of customer service" is related to the availability of home nurses for elderly persons living at home. These different "basic inputs" cannot be "summed" so it is necessary to complete each sub-domain analysis separately. However the conclusions are in the same form and, as already stated, can be combined to form the secondary general conclusions. In expert systems implementation terms the combination is best achieved by forward chaining inference. The resultant secondary conclusions are the important business performance indicators sought as the output for this type of task.

In Business Performance Assessment itself there is a further level of initial input data analysis not necessarily always found in this type of task. The initial input to the sub-domain analyses depends on a mechanistic calculation of a series of performance indicators on which sub-domain performance analysis is based. Such indicators were simple ratios such as those derived by dividing number of obstetric surgeons by number of midwives. These calculations are not relevant to the primary or secondary conclusions but are required to cancel out any effects due to size of business function considered in the analysis.

Multiple Simple Search with Secondary Grand Conclusion can be represented as shown on the next page.

BUSINESS PERFORMANCE ASSESSMENT
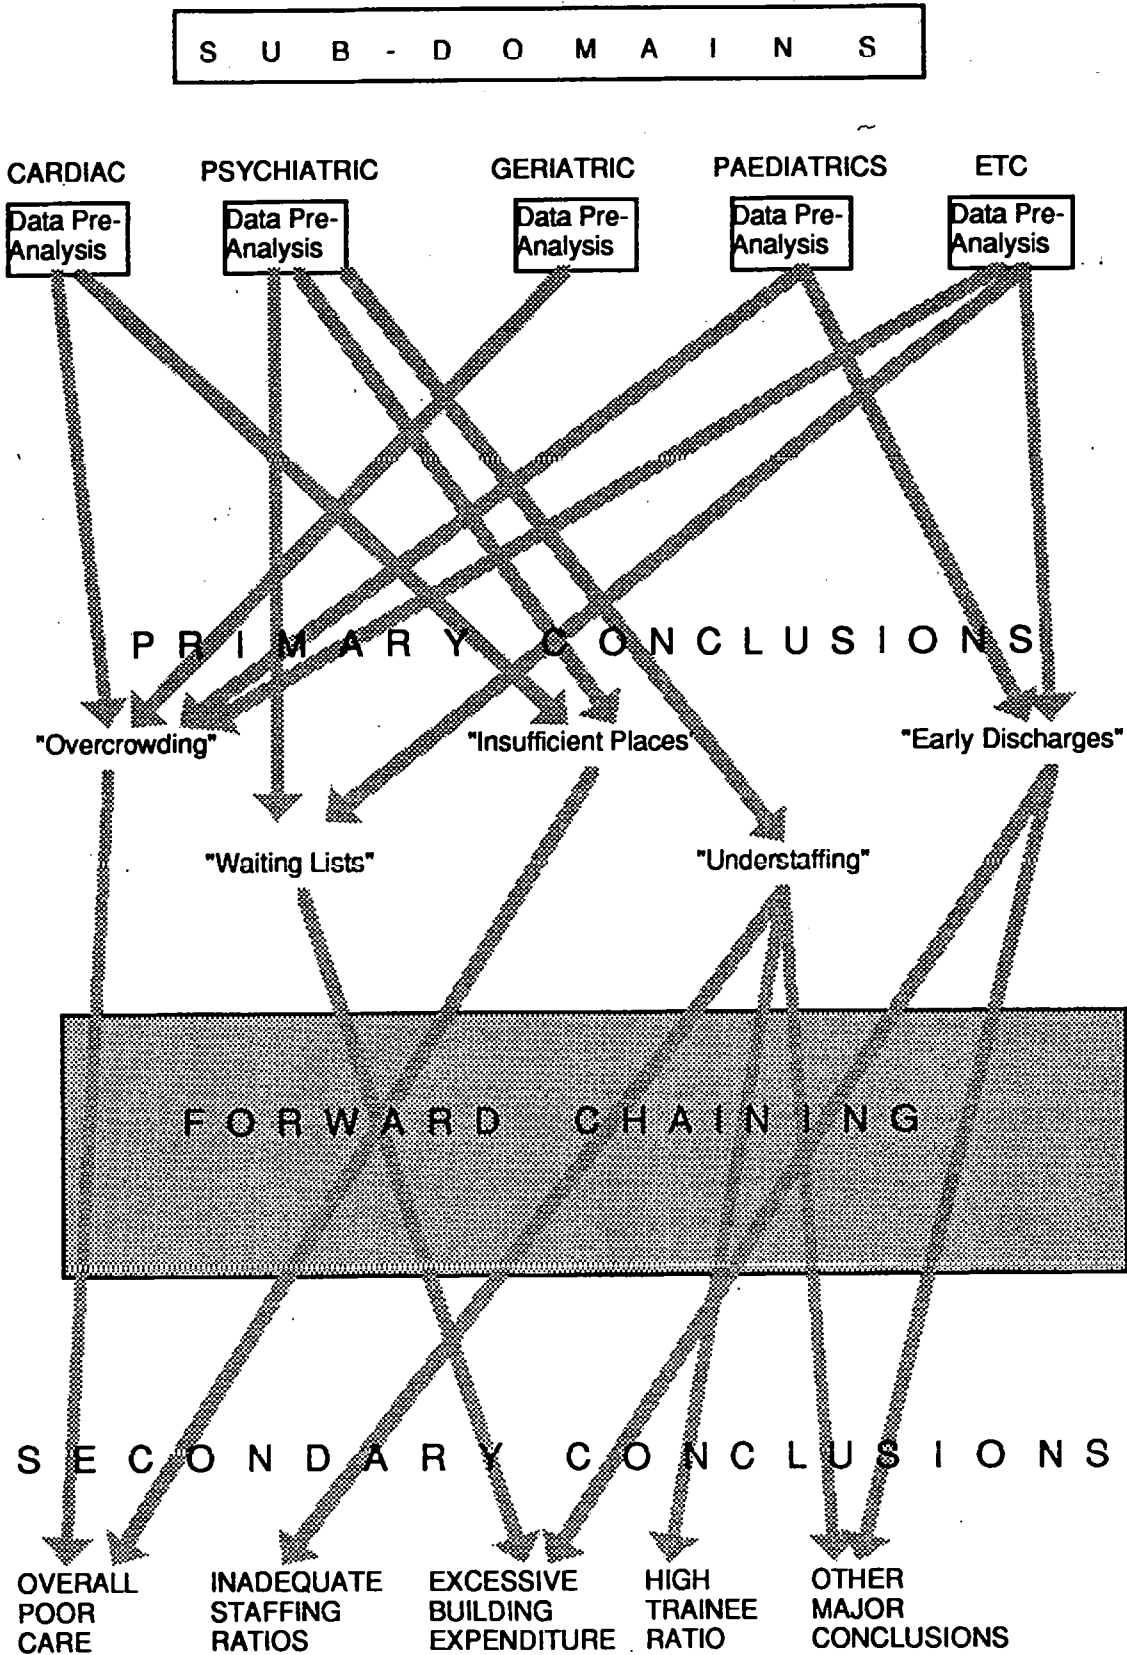
Diagrammatic Representation of Problem Solving Components

Business Performance Assessment has a similar representation to Audit Work Planning except that the Secondary Conclusion are termed "Grand" and are the primary result of problem solving rather than the "Limited" secondary Conclusions for Audit Work Planning. Further, once the analysis is completed, there is no further feedback or iteration as with Audit Work Planning.

Multiple Simple Search with Secondary Grand Conclusions

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Multiple Domains with
no Reference Linking

Grand
Overarching
Conclusions
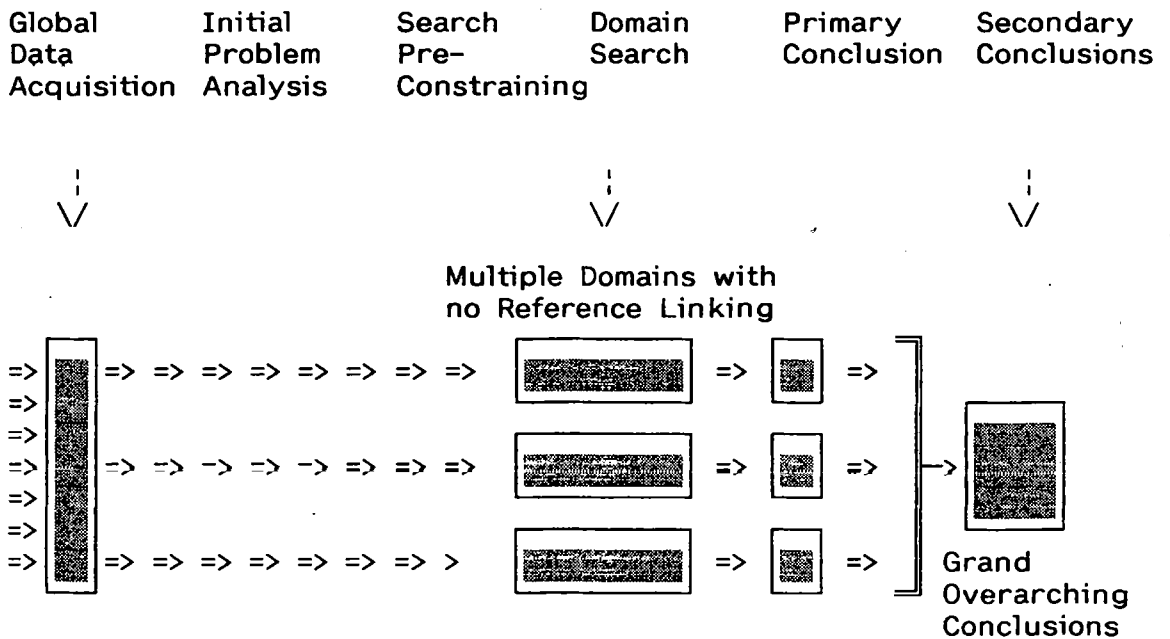
Commercial Implementation Details

The system was implemented on a standard PC so as to be widely usable at remote office sites. Crystal 3 was used to support the implementation with heavy usage of interfaces to Lotus 1-2-3 which was used to store the mass of performance indicator data.

The implementation was unusual in that, apart from user interaction to·select file data for analysis, set certain threshold values and name the file to hold the report, the system ran entirely in batch mode with no interaction from the user. Thus, once data for analysis had been selected, the system proceeded automatically to the point where the resultant report was copied to file.

This use of batch expert systems, and its extension to partially interactive systems, is considered further following this section.

Following deployment to the field, this system became really successful for reasons never anticipated in the original proposal or sponsorship.

In the time before hospital trusts etc, the managers of District Health Authorities were required to submit masses of performance indicator statistics (eg number of midwives per 1000 live births, number of elderly persons in residential care per 1000 population) to the Department of Health and these were laboriously analyzed so that an annual, including detailed recommendations to the Authority, could be prepared for the Minister on each Authority.

As already described, the system operated in batch mode and produced a well-formatted printed report. The quality of such reports was so good that some senior Department of Health staff were not aware that they had been prepared by the expert system. As with other systems described above (eg Audit Work Planning) presentation of output in a familiar form was very important for successful adoption by the user community.

Once the report had been considered by the Minister, and the Authority

manager criticised or congratulated, the report was sent back to the Authority where its advice was used to improve the Authority's performance.

Once the expert system had become established in use by the Health Authorities (partly as a result of publicity by the expert system shell providers) there was a demand by local managers to use the system locally rather than wait for up to a year for potential criticism and advice. As a result many copies of the system were sold and used on a regular basis by managers to "tune up" their performance rather than wait for the annual report to the Minister. The net result was that the advisory feedback loop to each Authority was made far more responsive both in terms of frequency of feedback and time to get advice on each occasion.

Thus, in summary, the purpose of the system changed from being solely a senior management annual performance monitoring aid to additionally being a "real-time" decision support system to the managers whose performance was initially being monitored.

Batch Processing Expert Systems

As explained above the implemented system was entirely data driven with no required interaction from the human user. Since the data from a number of Health Authorities was analyzed and reported in a batch, this system could therefore be classed as a "batch processing expert system". Many similar system can be envisaged. For instance, the author has a current interest in a major expert system which scans bank records looking for signs of multi-transaction fraud or other security problems. Another current batch processing expert system designs financial portfolios for

clients whose personal details and financial risk characteristics are recorded in a database.

In the case of the system described above, and possibly for other similar systems, there can be radical extensions of performance and "insightfulness" of the reports if some allowance for user interaction can be made.

For instance, not all problem data can be reduced to a database format. Thus, in the above task there was no way of encoding an operational unit's objective for growth (which might give an adverse effect on current year figures) or for "excusing" bad performance due to damage of buildings by floods or major staff absence due to epidemics. Similarly, some data may only required in a particular problem situation and to collect it on a routine basis would be grossly inefficient. Nevertheless, if such a problem situation is recognised, the system can request input of the extra data to complete an in-depth analysis.

In such cases there is a requirement for the expert system to take the analysis of the data to exhaustion and, where an exception is noted potentially for the purpose of the report, to conduct a diagnostic with the user to determine the cause of the exception and, where possible, offer useful advice for the future. In fact, in a later implementation of the above system, for which the author had very little personal involvement but for which he made design suggestions, this interactive element was incorporated with a real increase in utility of the overall expert system.

3.7   MULTINATIONAL TAX PLANNING

Sponsoring Organisation

The commercial investigation and implementation associated with the research on this task was sponsored by a major chemical company.

Introduction

The decision to pursue this general application area arose out of the author's wish to attempt to understand a representative area of complex senior managerial decision making. In commercial terms the investigation was an attempt to develop a saleable generic "high-level" financial expert systems application that could be profitably tailored to the needs of individual major corporate clients. In the event the specific sponsored project and expert system still in use are subject to major commercial confidentiality considerations. As a result the material reported in this thesis represents only part of the research outcome of a very worthwhile project and the author's desire to produce a generic tailorable system was not possible due to the success and resultant commercial confidentiality of the specific initial application.

The choice of the specific application area arose because of the availability both of commercial funding and first-rate domain expertise in multinational taxation. Other similar applications considered included, for instance, the commercial decision making involved in transfer pricing and marketing in multinationals.

Large multinational corporations consist of a group of companies, fully or partly owned, held in multi-stage hierarchical ownership relationships. It is usually the objective of the parent company to bring dividends back to the "home" country with minimal overall losses due to taxation.

The expertise of the corporate multinational taxation specialist, which was our primary interest here, is used to advise the corporate treasurer on how this optimum tax position might be achieved. The resultant commercial expert system implementation, MNCT, is a major commercial implementation of such expertise and provides a tool to assist the taxation specialist in his very complex task.

Author's Role

Apart from direct project management of the implementation team, the author's role was in both pilot and extended knowledge acquisition, analysis and MNCT system functional design. The actual implementation and system design, including much innovative user interface work was undertaken in Lisp on Symbolics hardware primarily by Malcolm West and Peter Richards. Both were then colleagues of the author when at Coopers & Lybrand. The implementation work is not directly reported in this thesis.

Outline Task Description

Multinational corporations evolve and grow over time. Such growth may be organic, where parts of the organisation expand as a result of business success or expansion of market demand. Alternatively growth may be by

acquisition or merger. At any one time the corporate tax specialist faces an existing corporate structure which, to a large extent, is the product of the historical growth of the organisation.

The tax specialist has an ongoing duty to find ways of optimising the overall tax burden to the corporation and also, on a "per occasion" basis, to recommend the tax optimum holding relationships of new acquisitions to the corporation.

Thus, in this latter role, the tax specialist is driven both by the ongoing tax minimisation goal and the specific circumstances triggered by the need to "slot in" a new company into the overall corporate structure.

Form of Investigation

Multinational taxation is a very specialised and complex domain well away from the author's own expertise. The experts in this area are usually very senior in their companies. Their level of personal expertise is probably higher than any of the other experts with whom the author has worked.

For this reason a longer than usual period of personal familiarisation was necessary before conducting the initial knowledge acquisition. This familiarisation took the form of a series of conversations with senior tax specialists together with study of books on multinational taxation.

Following familiarisation with the domain the first sessions with the corporation taxation expert were conducted. The taxation peculiarities of different multinationals vary widely and are very much a reflection of the

global structure and trading patterns of the organisation.

As a result the knowledge acquisition undertaken with the initial expert was heavily biased toward the specific characteristics of the organisation. Nevertheless, as a result of subsequent knowledge acquisition with two further tax specialists from other major multinationals it was possible to generalise the expertise involved in multinational tax optimisation. However, as reported above, commercial confidentiality considerations prevented this more general expertise being used for further applications.

The company specific tax information and strategies are obviously highly confidential to the companies involved. For instance, "tax holidays" may often be negotiated with the fiscal authorities in overseas countries and used exclusively by the company that has negotiated them. Widespread knowledge of such discretionary arrangements would attract similar multinationals to negotiate equivalent arrangements leading to a loss of commercial advantage of the initial company. Nevertheless the generalised skills are not confidential and show remarkable similarity between practitioners in different multinationals. It is these generalised methods of multinational taxation problem solving which are reported here.

Findings on Nature of Problem Solving

The tax expert is the custodian of the ongoing corporate structure and is called upon, perhaps several times per week, to suggest methods by which newly acquired companies may be assimilated into the corporate structure. He also has to devise specific plans for bringing back dividends to the UK in a way which avoids double or excessive taxation.

Problem solving commences with a thorough examination of the balance sheet and other financial schedules of the subsidiary (or potential subsidiary) company which is the focus of attention. In many cases the corporate tax specialist then consults his records to determine the specific current tax arrangements applicable to the country of residence of the focus company.

Often such enquiries are also followed up with verbal enquiries to counterparty tax specialists in the distant country or countries. The tax specialist then proceeds to construct a number of alternative diagrammatic relationships of companies associated with the target company and, sometimes with the aid of junior colleagues, to calculate the dividend implications of alternative holding and remittal arrangements.

By this route the problem solving is essentially focused to one small part of the organisation and, often with considerable difficulty, a proposed optimum structure is devised. This is then compared with the overall corporate holding structure to determine whether or not it is suitable from the global viewpoint. If not there are further iterative cycles until a solution which is both acceptable at the local and global level is found.

Problem Structuring Props (Tools and Aids)

The props used heavily in this task are the graphical trial partial corporate structures and the draft financial schedules. Such documents are prepared in a fixed format which, in the implemented system, is used to support a form of cooperative problem solving between the system and the human tax expert.

Generic Data

There is a mass of generic data associated with this task. Internally there are corporate specific data on the existing corporate structure, financial targets etc. Externally there is a requirement to assimilate an ongoing volume of published and distributed material on UK and overseas taxation rates and fiscal regulations. Some of this is immensely detailed and, for instance, includes data on which areas of which countries and for which periods certain commercial activities are tax exempt. For instance, certain areas of Ireland are from time to time given tax free status for certain industrial research activities. Similarly, certain regions of France are designated as low tax regions for specific industrial processes. Familiarity and knowledge of this body of international tax knowledge is one of the principal skills of the practitioner. It is also the reason why some experts concentrate on North American tax knowledge, some on European tax knowledge etc. Global tax knowledge at anything except a superficial level is impossible for any one human expert.

Specific Data

Specific data on target companies for acquisition or tax-optimised dividend remittal is provided by the balance sheet and other schedules, for current and historical years, of the target company.

Knowledge Sources

Two forms of knowledge are provided in the implemented system to

represent both group and company viewpoints on any suggested corporate structure.

Open/Closed Problem

The problem is an open problem where the solution can be any one of an infinite number of holding and dividend remittal relationships. Thus the chosen dividend structure can consist of any of a great many possible hierarchical ownership relationships, each of which can involve the paying of an infinite range of dividends (from zero to the maximum available from profit and accumulated reserves) to its parent.

Proactive/Reactive Problem

The problem solving specifically considered in the investigation (and the commercial implementation) is reactive to acquisition or the need to remit a particular dividend to the parent company.

Proactive problem solving is also undertaken as a background activity and as regular "housekeeping" exercises in order to maintain an overall tax-efficient corporate structure and dividend plan. Therefore the tax expert is regularly reviewing the overall corporate structure with a view to making it more tax efficient.

Problem Solving Architecture

In this task there is potentially an immense search space. For instance a

requirement to optimise the total corporate tax structure would result in a massive combinatorial explosion of possible holding relationships between companies. Thus, although the expert always stated that he had a duty to "optimise the overall tax structure of the group" which resulted in a degree of proactive exploratory activity (eg experimenting with the consequences of small changes in holding relationships), the task is mainly reactive to individual challenges and situations, such as mergers and acquisitions.

Such challenges always serve to localise the problem and so form a natural "focus of attention" constraining mechanism. For instance, the decision to purchase a Mexican company would serve to pre-constrain the task primarily to optimising the Mexican internal and external holding relationships. The potential external holding companies for a Mexican company would either be in the United Kingdom or Holland. (In the diagrammatic representation of problem solving components later in this section we refer to this as "institutional pre-constraining" of the search since it is a deliberate company policy to optimise locally and then, as described below, to feed the results back for reconciliation with the overall group tax situation.)

As a result the overall "problem" would be constrained to Mexico (including internal companies) and United Kingdom and Dutch holding companies. Even within this set of companies there are further institutional constraints due to "stiffness" of historical holding relationships which cannot be changed except under exceptional circumstances. (Similarly, in a parallel international tax planning case, investigated for another major multinational but not yet commercially implemented, it was impossible to re-order the holding relationships of certain Italian subsidiaries without agreement of the Italian Mafia. Such permission was usually difficult to obtain.)

The search for the best structure within the limited set of companies and restricted possible relationships is heuristically guided but, in the extreme case, can be undertaken almost on an exhaustive search basis. Once the constrained problem has been "solved" the tax implications of the locally derived solution are reflected back to determine their impact on the tax situation of the overall group.

If there is a particularly poor overall implication from this "feedback" the result is taken into consideration during another iterative pass through the problem.

Multinational Tax Planning can be usefully compared with Permanent Health Insurance considered later. In the latter case, but for somewhat different structural reasons, there is also the likelihood of massive combinatorial explosion. In Permanent Health Insurance there are no constraining measures and problem solving is often non-optimal. In contrast Multinational Tax Planning provided an outstanding and industry-respected example of expertise in this difficult field.

Multinational Tax Planning is illustrated on the next page:

MULTINATIONAL TAX PLANNING

SPECIFIC
TAXATION
CHALLENGE

Europe   Asia   Americas   Australasia   CONTINENT

UK F NL etc   Ch J HK etc   US C Mx etc   Aus NZ etc   COUNTRY

New Constraint → UK   NL   Mx   CONSTRAINED PROBLEM

NEW HOLDING OR DIVIDEND STRUCTURE

Unfavourable Impact

IMPACT ON OVERALL CORPORATE TAX SITUATION

ACCEPTABLE IMPACT

FINISH OF CONSTRAINED PROBLEM

104
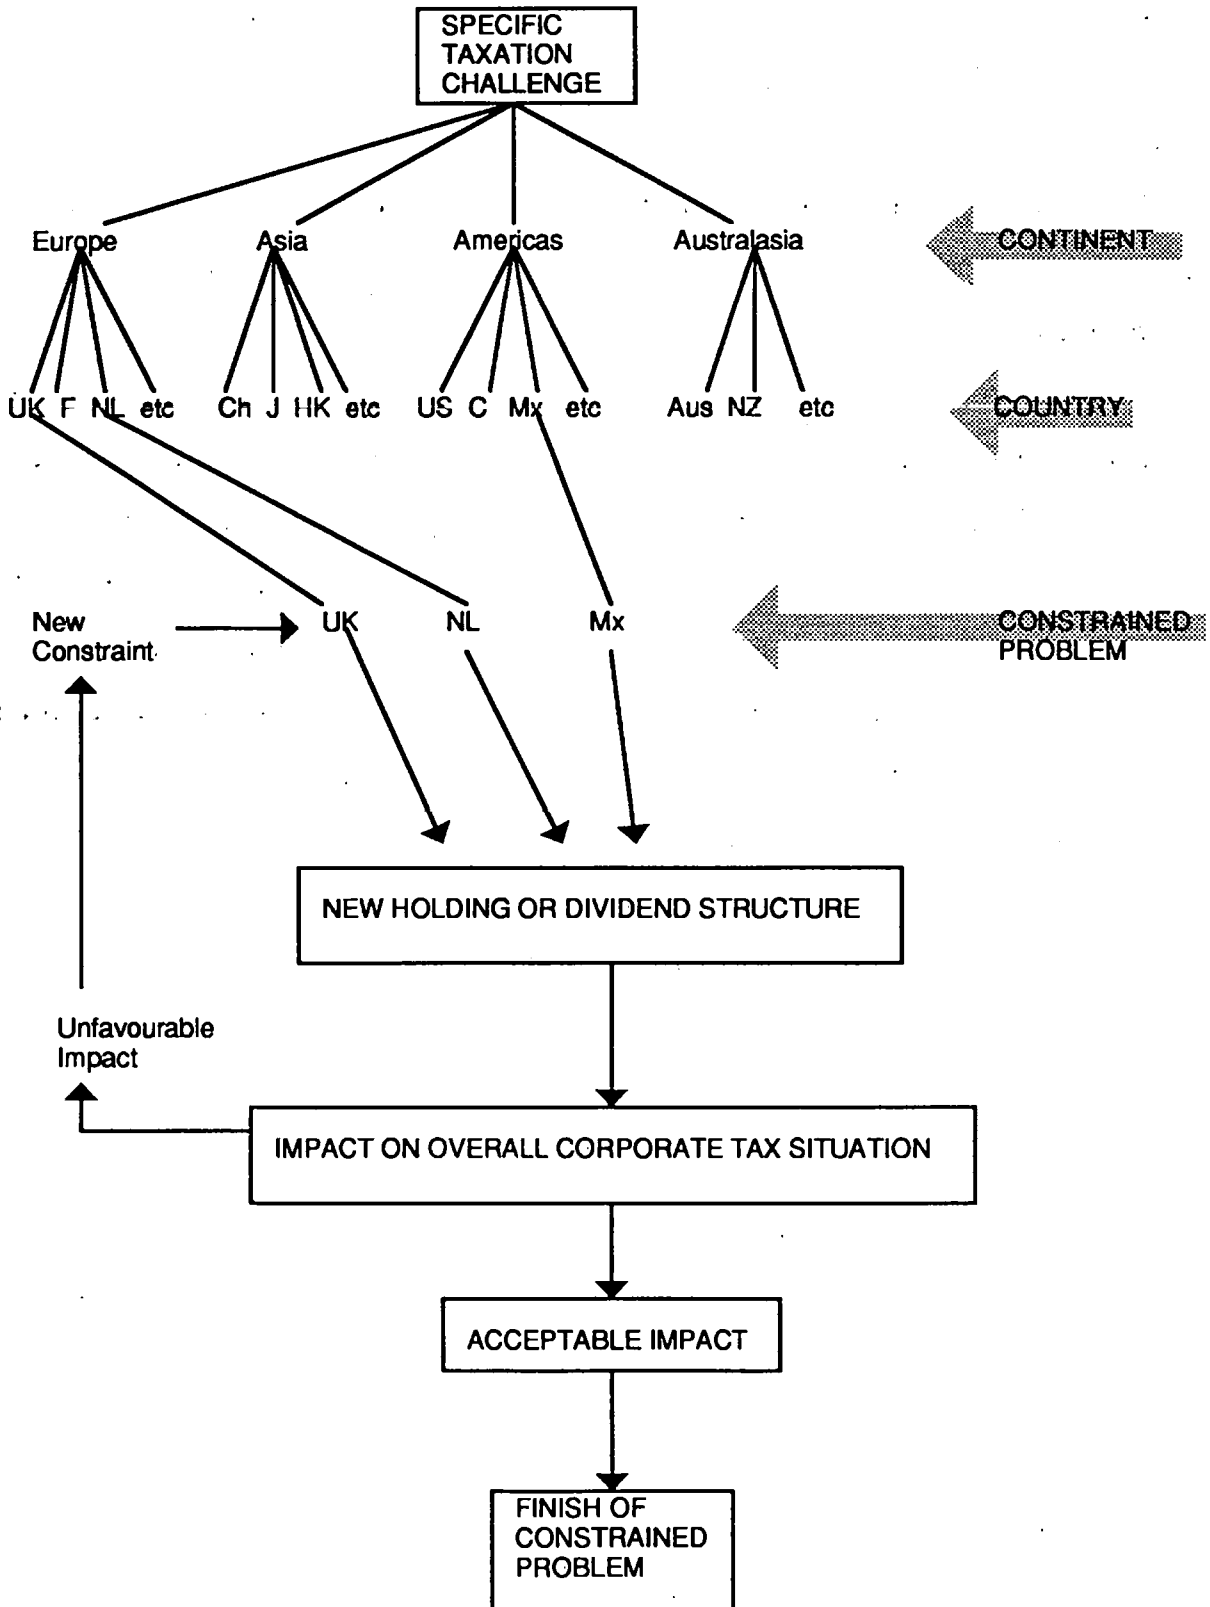
Diagrammatic Representation of Problem Solving Components

In the diagram it can be seen that an initial component of problem solving constrains the search to a single domain (eg that concerning the tax arrangements for Central and South American companies). This constraint is based on the practicalities of the institutional problem solving situation. (In comparison, the problem solving in Engineering Preventative Maintenance which follows is constrained by the nature of the domain.) The searched domain is shown as the darker shaded lower domain on the diagram. Once a conclusion to the constrained problem is reached then the new situation is assessed in the perspective of the overall or global planning tax arrangements for the group. If the local solution is at major variance with the overall plan then the tax expert iterates through the problem with the additional information as just one additional input to his constrained problem solving. For instance, a decision to parent a newly purchased subsidiary on a Mexican holding company could severely disrupt a global plan to delay making a dividend in Mexico and, in view of the overall situation, it may be better to parent the new company elsewhere.

Pre-Constrained Search Tree with Local Result Affecting Overall Result

| Global Data Acquisition | Initial Problem Analysis | Search Pre-Constraining (Institutional) | Domain Search | Primary Conclusion | Secondary Conclusions |

Constrained Domains with no Reference Linking

Global Feedback
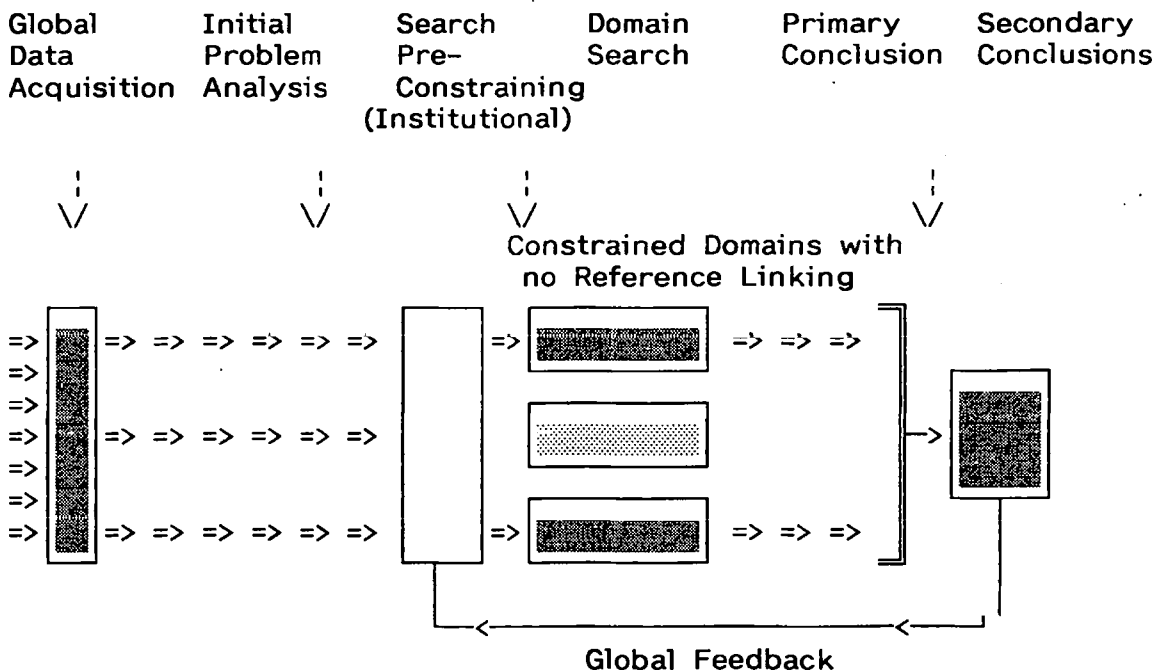
Commercial Implementation Details

This system was extremely difficult and time-consuming to sell to the sponsor yet is now in very confidential, successful and profitable use.

The tax expert was very supportive but the nature of the implementation demanded that it be delivered on Symbolics hardware. The company had established a tight mainstream computing strategy (IBM mainframe etc) and considerable time was spent justifying the need for Symbolics hardware. Similarly interfacing to the mainframe database was an issue that had to be solved in advance of system authorization and implementation. In order to persuade the customer as to the potential benefits of the full MNCT system a pilot system, CLINTE (Gleeson & West, 1988), was implemented initially using Inference ART on a Symbolics 3670. Considerable use was made of bespoke graphics written entirely in Lisp. As a result of the success of the pilot system the full MNCT system was implemented using the Symbolics Joshua toolkit.

The expert system itself was complex but the main problems concerned the user interface. Tax experts think in terms of company balance sheets and schedules and it was imperative that the system could display such schedules in exactly the same format as would be used on paper. In fact extra facilities were offered in that a sort of hypertext functionality was offered so that, by clicking on a particular balance sheet figure, the figures and trends for related years would appear together with subsidiary data on how the figure was generated. Thus the normal two-dimensional schedules were given a form of three-dimensional functionality with the ability to "plunge down" on any figure to discover both the previous years' figures "underneath it" plus the data on the subsidiary companies on which the

figure was based. Therefore, as an adjunct to the main tax planning system, the user was offered a means of exploring the figures on which the system based its reasoning.

The system output also had to be in a form very familiar to the user. Organisation charts for multinational companies usually show an hierarchy of several hundred companies together with the tax holding and mixing arrangements. Tax planners use such charts as a problem solving prop when trying to devise new holding arrangements. Therefore the system had to display and print such charts as its output. The limitations of the A4 printer, which produced charts in sections, was one of the unresolved negative features of the system.

The intended user was to be an existing expert but, most often with tax knowledge on a different part of the world. For instance, for the reasons explained above, the system might be used to give advice on European tax holding relationships for a user who was expert on North American tax arrangements and legal requirements. In such cases the user will want to improve the advice suggested by the system to explore even better tax relationships based on his own knowledge. In the main these enhancements will be tuning of holding relationships based on an overall tax structure suggested by the system. The net result is that the user required the system to have a graphical based "what-if" system. Therefore, by dragging a subsidiary into a new tax parenting relationship the user would want to see the overall effect on revenue flow back to the parent company. This facility had nothing to do with the core expert system but entailed considerable effort to make the overall system effective and acceptable by the users.

Additional Comments

It is still hoped that this investigation of multinational taxation skills will eventually lead to several commercial implementations and, subject to the resolution of commercial confidentiality issues, the possibility of a generic tool for European tax applications.

3.8   ENGINEERING PREVENTATIVE MAINTENANCE

Sponsoring Organisation

The investigation was sponsored by a major cigarette manufacturer faced with the problem of increasing the cigarette making machine utilisation of a particular plant manufacturing an increasingly successful low tar cigarette.

Introduction

The task described in this section, Engineering Preventative Maintenance, was investigated a considerable period after the other major engineering tasks, Benchtop Electronic Repair and Telecom Repair Service. Although basically a diagnostic task, there is a marked difference between the mechanical domain and the electrical/electronic domains represented by Telecom Repair Service and Benchtop Electronic Diagnosis. In electrical domains the only means to verify a hypothesis is by replacement of a component or other physical circuit intervention. Mechanical diagnosis is different in that, once the equipment has been disassembled, the fault can usually be detected and verified visually.

However, although the visual dimension makes mechanical diagnosis easier, there is still a price to be paid for disassembly, replacement and adjustment. Unsoldering a potentially faulty component may be a nuisance but it is far quicker than unbolting and removing several "layers" of heavy mechanical assemblies to check and replace a particular component. (The

situation may be better understood by analogy to car mechanics. A problem in engaging first gear may be due to a worn synchro-mesh gear which can only be accessed and replaced by first unbolting and removing the entire gearbox from the car, then removing the relevant shaft from the disassembled gearbox so that, finally, a new gear may be replaced on the shaft.)

Mechanical systems also have the problem of wear. In general a piece of electronic equipment either "works" or "fails". Mechanical systems gradually wear and degrade and, unlike electronic components, mechanical components have a finite life after which they must be replaced. Such issues make mechanical systems maintenance very different from electrical and electronic fault diagnosis. In fact, it seems incorrect to speak of "electronics systems maintenance" since there is really no equivalent to the preventative or routine maintenance undertaken on mechanical systems.

Author's Role

The author was solely involved in this activity, undertaking initial investigations, knowledge acquisition and analysis.

Outline Task Description

The task expert was the maintenance manager of a large cigarette factory. Cigarettes are produced on complex three-stage machines and these large, primarily mechanical, systems require both routine and breakdown maintenance. The manager's task is to optimise both forms of maintenance

activity so as to give maximum available machine utilisation and so, indirectly, maximum overall cigarette output.

Form of Investigation

The factory layout and systems were well-known to colleagues of the author and these persons were informally interviewed in advance to gain an understanding of the complex cigarette making process and its problems. This was followed by a detailed guided tour of the factory floor by the expert including an opportunity to examine and understand the mechanics of a partly disassembled cigarette making machine. Finally, knowledge acquisition was undertaken involving the exhaustive examination of machine maintenance documentation and the expert's notes and schedules.

Findings on Nature of Problem Solving

The expert was a particularly conscientious and methodical worker who had deliberately studied his problem area over a number of years in the belief that he could achieve much higher machine output than claimed by the machine manufacturers. As a result he had demonstrated that he could achieve significantly better optimised maintenance (and so cigarette output) than the shop floor staff normally responsible for routine maintenance planning. The objective of the proposed expert system was to deliver the manager's expertise to shop floor staff and so increase machine utilisation from about 60% to a more acceptable target figure of 80%.

There were two aspects to the problem solving. First, maintenance had to

be planned on a routine basis to replace components subject to normal mechanical wear and abrasion.

Since components wear out at different rates, scheduled replacement cycles vary widely from component to component. For instance, apart from lubrication, rollers in the case making assembly require replacement on a yearly basis. At the other extreme, nylon belts delivering tobacco into the paper normally require replacement on a daily basis.

Apart from the planning of routine maintenance, it was necessary to respond to mechanical failures with breakdown maintenance. At such time it was necessary to effect a diagnosis (usually a fairly straightforward task with mechanical systems) and, using the opportunity provided by the necessary disassembly of the machine, both make a repair/replacement and effect any routine maintenance that could be conveniently undertaken.

The expert's skill was in optimising routine maintenance (so as to increase utilisation without undue risk of sudden failure) and to make the best and most productive use of the disassembly opportunities resulting from a breakdown.

Problem Structuring Props (Tools and Aids)

The expert used both the staff activity work schedules (a type of diary of planned work for each individual maintenance worker) and the maintenance record logs of individual machines to devise and experiment on optimum routine maintenance schedules. When a breakdown occurred these documents were used on an immediate basis to allocate maintenance effort

112

to the breakdown, schedule opportunistic maintenance and re-plan the staff/machine maintenance schedules disrupted as a result of the breakdown.

Generic Data

The cigarette making machines (from two different manufacturers) were provided new with comprehensive maintenance schedules for all mechanical parts. (These are similar to a car maintenance handbook advising changing break pads every 5,000 miles, new spark plugs every 10,000 miles etc.) Thus manufacturers recommended tobacco cutters should be replaced every 1000 hours and paper grippers every 500 hours. As a result of his experience of seven years maintenance of the machines, the expert had 'tuned' this generic data so that, for instance, cutters were optimally replaced every 1200 hours and grippers every 1000 hours.

Specific Data

Although the overall maintenance is based on the schedules just mentioned, individual machine records form the specific data on which is planned routine maintenance is based. These records are also used at the time of breakdowns to determine, for instance, which routine maintenance is nearly due and can be efficiently brought forward as a result of the necessary stripping-down of the machine.

For instance, a shaft may be close to the site of a worn component being replaced as a result of a breakdown. If the machine records indicate that the shaft is within, say, 100 hours of a routine 1000 hour maintenance

replacement then that activity may be brought forward to the time of the breakdown maintenance. The individual machine record would then be updated to form specific data on which to base the next routine maintenance cycle.

Knowledge Sources

The initial implementation was planned to deal only with the management of breakdown repair. The planning of routine maintenance is more straightforward and, for the purposes of this investigation, not so interesting.

Two separate knowledge sources were envisaged for the breakdown maintenance implementation. The primary knowledge source facilitates an essentially backward reasoning search in a search tree pre-constrained by the nature of the mechanical site of the fault. The second knowledge source is forward reasoning and fragmented in the same way as the sub-search areas of the primary knowledge source. The second knowledge source acts both upon the requirements to disassemble parts of the machine (necessary for the primary diagnosis/repair/replacement) and the current specific maintenance state of related components to plan opportunistic maintenance replacement for those nearby components.

In a more complete implementation routine maintenance would be supported procedurally. However, and not fully investigated, there would be an opportunity to offer some form of inductive refinement system to undertake the ongoing optimisation and tuning of the generic schedules discussed above. Thus routine replacements of a component found regularly not to be

worn would result in a lengthening of the replacement schedule. Similarly, regular breakdown before the scheduled maintenance time would result in a shortening of the schedule.

Open/Closed Problem

The problem is closed to the set of replaceable components that constitute the machine. In some cases components can be either worn, broken or maladjusted. In most cases the only valid symptom is wear causing the machine to malfunction.

Proactive/Reactive Problem

Reactive to breakdown with proactive opportunistic maintenance at the site of the repair/replacement.

Problem Solving Architecture

This type of task is common with Multinational Tax Planning in that there is a potentially very large search space which is pre-constrained prior to the start of the analysis of the problem. However the use of the result of the constrained search is rather different to Multinational Tax Planning.

Also Engineering Preventative Maintenance involves diagnosis but differs from Benchtop Electronic Diagnosis in several ways. In Benchtop Electronic Diagnosis the search tree may be dynamically re-arranged as a result of

information in the initial problem description (ie the Fault Docket). However, at the extreme or in cases of ill-described faults, the whole search tree is available for search.

In contrast, in mechanical diagnosis problems such as Engineering Preventative Maintenance there can be positive visual or logical pre-constraining of the search due to the isolation of mechanical sub-systems. Thus, in Engineering Preventative Maintenance, a fault in the Case Making part of the machine can have no cause in the Rod Maker. In many cases the constraining of the problem can be to quite low-level mechanical sub-assemblies even before any physical diagnosis is undertaken.

Thus there is some real commonality with Multinational Tax Planning although the particular example of Multinational Tax Planning has a constraint which is pragmatically or operationally imposed rather than imposed by the nature of the domain. However in Multinational Tax Planning there is essentially a "feedback" loop from the problem solution to the initial (constrained) problem description.

In Engineering Preventative Maintenance is no such feedback since the solution to the local problem can have no impact on the overall maintenance considerations. However, local fault repair maintenance activity can give rise to further locally related opportunities for preventative maintenance.

Thus in this task the activities necessary to undertake the fault maintenance are used as inputs to localised opportunistic reasoning (implemented as forward chaining and considered as a distributed but discrete supplementary knowledge source), to discover cost-effective preventative maintenance that may be undertaken concurrent with the fault

116

maintenance.

In summary there is pre-constrained search but the results of the localised search are used to optimise related local activity rather than the global situation. Engineering Preventative Maintenance is illustrated on the next page.

ENGINEERING PREVENTATIVE MAINTENANCE

```
                    ┌─────────────┐
                    │  SPECIFIC   │
                    │FAULT REPAIR │
                    │MAINTENANCE  │
                    │  PROBLEM    │
                    └─────────────┘
```

Rod
Maker

Maker
Assembly

Packer

SUB -
EQUIPMENTS

Tape
Feed

Rod
Former

Spider

Cutter

Assembler

Box

Case

SUB-
ASSEMBLIES

Paper
Feed

Gripper
Assembly

End
Stop

Rotator
Assembly

Bander

IR-Inspect

Roll
Assembly

Glue
Feed

Filter

Filter
Band

Cutter

Turner

ACTION CONSTRAINED TO
GRIPPER ASSEMBLY OF SPIDER

SPECIFIC FAULT: Gripper Collet

OPPORTUNISTIC FORWARD CHAINING::
Replace Gripper Feed Fingers
Replace Spider Drive Cog Chain

SPECIFIC
FAULT
REPAIR

PLUS

OPPORTUNISTIC
PREVENTATIVE
MAINTENANCE

Diagrammatic Representation of Problem Solving Components

As with Multinational Tax Planning the search is constrained to the a particular sub-domain, that of the Gripper Assembly. However, this pre-constraining is a consequence of the domain. Whereas in electronics a fault in one area of a circuit has some chance of causing a problem in a remote circuit element, or in International Tax Planning a tax decision in Australasia could affect the overall group, there is no way that cigarettes could be produced in unequal lengths by the part of the machine that puts sealed cartons into cardboard packing cases. Therefore, the search is "domain constrained" to a single sub-domain.

The objective of the problem solving is to determine the faulty component as a primary conclusion to the search. Local opportunistic problem solving takes advantage of the disassembled machine, and the visual clues that this presents, to "forward chain" (using the expert systems terminology for the human problem solver) to conclusions regarding replacement of components that may fail in the near future.

The Local Opportunistic Conclusions component is not shown in any of the other tasks reported. However, it is quite common in "mechanical" diagnostic tasks. For instance, in the gynaecological laparoscopic investigation of Fallopian tubal patency the surgeon will usually take the opportunity of, for instance, looking for perineal adhesions of the bowel and other organs. If these can be treated while the patient is still "disassembled" then, although non-symptomatic at the time of surgery, problems in the future may be avoided and further hospitalisation may not be necessary.

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining (Domain Driven)	Domain Search	Primary Conclusion	Secondary Conclusions

Pre-Constrained
Single Domain

=> => => => => => => =>
=>
=>
=> => => => => => => =>
·=>
=>
=>·=> => => => => => =>

->

Local
Opportunistic
Conclusions
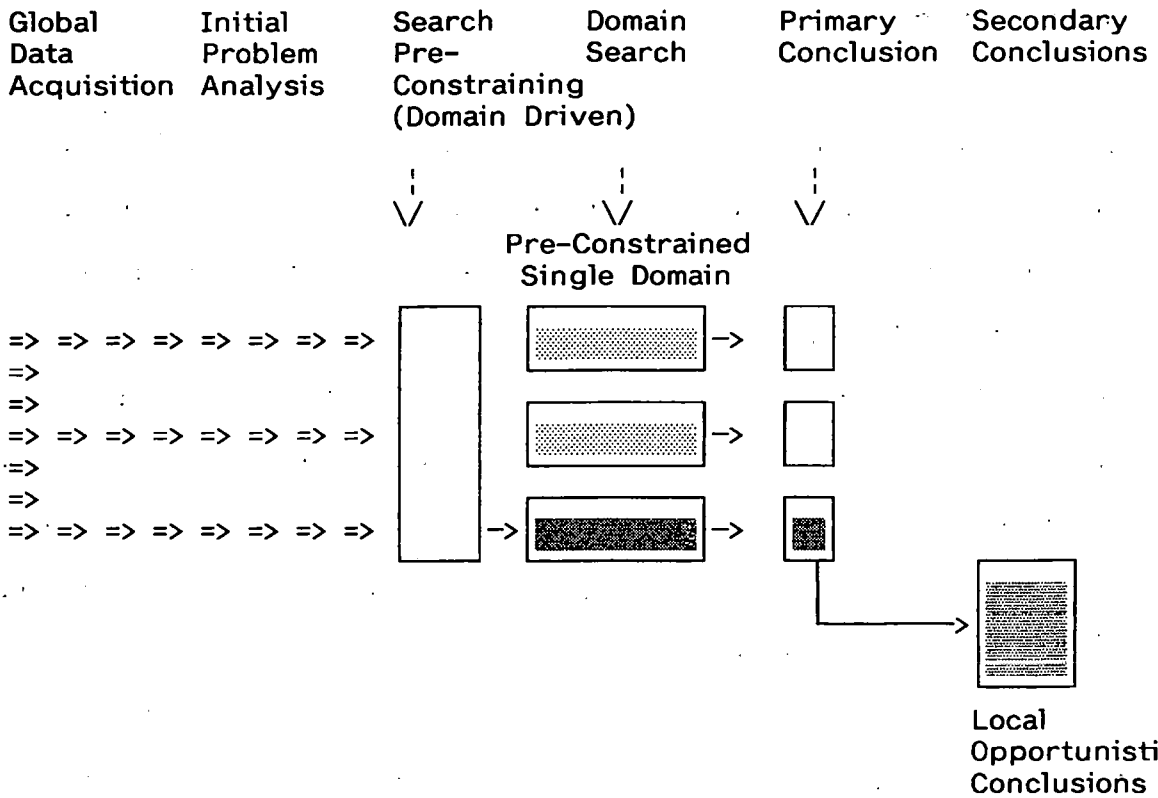
Commercial Implementation Details

Very unfortunately, the Engineering Preventative Maintenance investigation did not result in a commercially implemented system. Nevertheless, the investigation proved a very worthwhile experience for the author. There was considerable knowledge involved in the task and the commercial implications were very extensive. Well chosen maintenance interventions, particularly the optimisation of opportunistic maintenance, could increase overall plant utilisation by several percentage points. Factory productivity and the return on plant investment could thus be radically improved.

Although this system was never implemented the author is convinced that a first rate system could have been constructed, largely as a result of the extent and depth of expertise of the human expert. The human expert, supported by a mass of wear data which he had painstakingly assembled, could accurately guess when maintenance should be undertaken and thus

prevented many daytime breakdowns.

The problem which prevented an implementation being undertaken was simple. The cigarette company management, although admitting that the expert was a "good guy", just did not understand what he did or how valuable was his expertise. In fact it was quite difficult to isolate him under the layers of very polite managers in the factory. Frankly speaking, the man's rank (he was graded as an assistant foreman) did not match his actual usefulness to the company and nobody would acknowledge that a worthwhile system could, or ought, to be built using his knowledge. Higher level staff were preoccupied with elaborate "top down" maintenance plans which they considered should be the basis of any system to be implemented and were unimpressed with the bottom up "craft knowledge" of the real expert. In fact the expert only showed his notebooks and accumulated wear data to the author after considerable trust had been established. Up to that point this material had been entirely ignored by management.

In summary, an excellent preventative maintenance expert system was never built purely because the expert was insufficiently cherished by his management. It was not the purpose of this research to investigate the merits, rewards and status of individuals in organisations. It is unlikely that the author, whose business status was that of a systems rather than an organisational consultant, could have affected the entrenched understanding of the status of the expert in this case. However, as a side conclusion to this investigation we must note that this case highlights the potential serious commercial costs and consequences resulting from situations where there is a managerial lack of appreciation of the skills held by more junior members of the workforce.

Additional Comments

This task is very typical of many mechanical engineering maintenance tasks. Diagnostic problem solving is very straightforward in contrast to electronics maintenance problem solving where visual diagnosis cannot be applied.

Nevertheless mechanical wear necessitates regular replacements (unlike electronics where "wear" is unpredictable and catastrophic) and there is considerable scope for systems optimising routine and breakdown maintenance.

3.9   PERMANENT HEALTH INSURANCE

Sponsoring Organisation

This initial investigation was at the invitation of a UK insurance company. The production of quotations for permanent health insurance was believed to require considerable real expertise and an attempt to support this with an expert system was considered very worthwhile.

Introduction

Most insurance related to continued personal good health is life insurance, ie insurance against death from natural or accidental causes.   However, high-earners and vital corporate employees are often subject to a far less common form of insurance, namely Permanent Health Insurance.   This is an insurance primarily to protect an individual against being unable to follow his or her existing occupation as a result of unexpected health problems.

Whereas life insurance, as described in the following section on Life Underwriting, is arranged on the basis of expediency with considerable "averaging" of risks over an immense number of policy holders, permanent health insurance is arranged on the basis of an accurate assessment of the risks involved.

This report covers only the initial task investigation and outline knowledge acquisition.   This material is entirely adequate for the purposes of this thesis and a good understanding of the expert's problem solving was

acquired. Nevertheless, the key expert involved was particularly resistant to computer support systems and it was not possible to further the investigation toward what could have been the basis of a very powerful and useful expert system.

Author's Role

The author was solely responsible for making this investigation and for the analysis of the acquired knowledge.

Outline Task Description

The expert task is to analyze the risk of incapacitating health problems on a named individual and prepare a quotation for cover based on the likely impact on employability and income.

Following receipt of a very detailed proposal form there is in nearly all cases a full medical examination of the applicant (often including X-rays, blood tests etc) and this is used by the expert as a partial basis of his quotation.

Form of Investigation

The task was introduced to the author as being of considerable importance but confined to a very small number of expert underwriters in the company. Before this investigation was undertaken senior management believed that

the underwriter's key expertise was in the area of interpretation of complex medical reports. As discussed below, this belief is probably false.

A short briefing session with a subordinate of the key expert was arranged but, in the main, all the knowledge had to be acquired from the rather reluctant expert himself. However, although very negative to computer systems, he was very co-operative during the protracted single knowledge acquisition session and, together with follow-up access by telephone, most of the available core strategic problem solving knowledge was acquired.

Findings on Nature of Problem Solving

As stated above, it was thought that the primary aspect of problem solving was interpreting the level of seriousness of a declared (ie on proposal form) or revealed (ie by medical examination) medical health condition.

However, although this was important, the health aspect was one of at least four "input parameters" which required quantification as inputs to a complex combinatorial evaluation upon which the premium was directly based.

The other primary input to this combinatorial evaluation was the level of risk due to occupational, recreational or residential factors. Thus a mining engineer, pot-holer or resident of the tropics could each be considered to have higher than normal environmental risks. (In the first two cases such risks are purely due to the activities undertaken; in the third case the extra risk is partly due to the tropical environment (eg malaria risks), possibly due to terrorism or war in certain countries and partly due to the potentially sub-standard health care available in the case of illness or

injury.)

.The third input factor was an assessment of the impact of disability on the applicant's ability to work. Sedentary workers (eg computer programmers) can tolerate a much higher level of disability than manual or mobile workers. Thus paraplegia would effectively end the career of a oil rig engineer but might have a much lesser effect on an architect.

The final, and most straightforward, factor contributing to the decision concerned the level of coverage required. Thus the policy might allow for full replacement of anticipated earnings or some proportionately lesser sum just reflecting the underlying incapacity due to ill health. Nevertheless, even in this case, the policy might allow for replacement of earnings and this would require some judgement of the future earnings potential of the applicant.

The most important and difficult aspect of the decision making was in the combination of the results of the input parameter evaluations to make an assessment of overall risk. Because of the wide range of combinations of the four input parameters, with the prospect of many "novel" cases, there was often a need for reasoning from first principles with little reliance on experience from previous cases. This aspect is considered further in the sections below.

Problem Structuring Props (Tools and Aids)

As with Commercial Business Underwriting the principle structuring tool was implicit in the layout of the basis of quotation on plain paper. Nevertheless,

126

the strength of this virtual tool, even though it was useful, seemed far less important than in Commercial Business Underwriting and, indeed, much of the reasoning was not committed to paper.

Generic Data

The record of previous quotations provided a library upon which decision making on a current case could be based. Thus cases which had demonstrated special medical or environmental risks and the expert stated that these were drawn upon if these were encountered again. However, there was no proper system for accessing these cases and, when it was done, it relied upon the expert's memory and was not reliable.

To a limited extent special rates tables were used for rating uncomplicated single risk factors. Thus a straightforward "excess" could be applied to a stable diabetic provided that there were no complicating occupational or other factors. However, such special rates tables were not effective in combinatorial risks where, in the worst case, an established health problem was likely to be aggravated by a particular occupational risk and that the resultant disability would be incompatible with continued work.

Specific Data

The detailed proposal application forms and medical reports, annotated where necessary by the insurance company's medical practitioner, were the sole sources of case specific data.

Knowledge Sources

Discrete sources of knowledge to quantify the health, environmental (ie occupational, recreational and residence) and occupational consequential risks were employed. A further, less extensive, source of knowledge is required to evaluate level of coverage. The final major source of knowledge is that required to deal with the reasoning required to combine and evaluate (or recognise) the primary risk assessments and so prepare a quotation.

Open/Closed Problem

In theory the problem is closed but the very large combinatorial matrix to cover the possibilities arising from the primary risk assessments makes it, in practice, an open problem.

Proactive/Reactive Problem

The task is solely reactive to the invitation to quote.

Problem Solving Architecture

Permanent Health Insurance has an initial resemblance to Business Performance Assessment. Business Performance Assessment is characterised by Multiple Simple Search with Secondary Grand Conclusion. However the conclusion element of Permanent Health Insurance is truly combinatorial with many thousands of possible combinations derived from the results of the

primary searches.

In Permanent Health Insurance the primary searches are shallow so, for instance, the "Age" search has only a depth of one level and merely classifies the applicant into age bands. Similarly the "Occupation" search has a maximum of three levels of depth and most occupations are classified at a single level. However the search tree is very wide. Similarly the "Hobbies and Sports" tree has a maximum of two layers in depth (eg "Rugby Union" and "Rugby League" are dependent on "Rugby", "Regular" and "Holiday" are dependent on "Skiing") but is also very wide.

The number of combinations is seemingly the product of the numbers of terminal nodes in the seven sub-domain simple searches. This can be considered as a seven-dimensional matrix where only a minority of the cells have known values. In this task it appears that novel combinations are either refused, dealt with if closely analogous to a recognised known cell or dealt with from first principles. This scheme of problem solving is laborious and often gives rise to idiosyncratic results.

Permanent Health Insurance is illustrated on the next page:

PERMANENT HEALTH INSURANCE

PERMANANENT HEALTH
INSURANCE GOAL

AGE SEX   MARITAL   HEALTH   OCCUPATIONAL   SPORT   OCCUPATIONAL   COVER
          STATUS   RISK   RESIDENCE   NEED   %
                                            HOBBY RISK

M F   S M D   Manual Clerical etc   Manual Clerical etc

<20 20-30 etc   Nil Chest Cancer etc   Soccer Iran Gliding etc

MATRIX:
Age x Sex x Marital x Health x Occupational Risk x Sport etc x Occupational Need x Cover

KNOWN CELLS   ANALOGOUS CELLS   UNKNOWN CELLS

Easy Quotation   Easy Quotation   Refusal   "First Principles"
(Acurate)   (Doubtful)   (Failure)   (Very Doubtful)

Diagrammatic Representation of Problem Solving Components

This task is straightforward to represent as components. As seen before,
there is global data acquisition (most of which is encoded on the application
form) following which the underwriter considers each of the sub-domains.
For instance, he will reach separate conclusions on the applicants's health,
working risks, work requirements, likely future earnings etc. Once these
primary elements of a decision are assembled, the underwriter will attempt
to combine them all into a single conclusion resulting in a premium.
Although the individual sub-domain problem solving is elegant, it is this
latter "massive combinatorial" part of the decision making that we feel is not
very well principled.

Multiple Wide Shallow Search with Massive Combinatorial Grand Conclusion

Global Data Acquisition	Initial Problem Analysis	Search Pre- Constraining	Domain Search	Primary Conclusions	Secondary Conclusions

Multiple Domains with no Reference Linking

Massive Combinatorial Conclusion
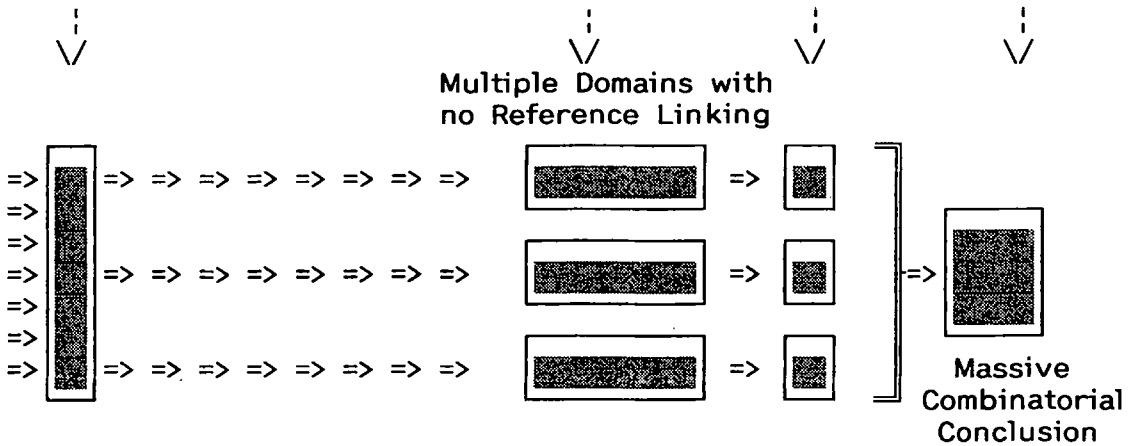
Commercial Implementation Details

Although a tentative implementation design was proposed in advance of detailed knowledge acquisition the author is not convinced that an expert system could ever be implemented for this application. However, if further investigation showed that such a system were actually feasible, it would appear that an expert system in this area could both be very cost-effective and radically improve accuracy of quotation. These effects would be particularly apparent as the system matured and the knowledge base expanded to encompass principled evaluations of the wide range of risk combinations involved.

In the particular case here reported the expert was rather uncooperative and full of his own self-importance. However, he was very authoritative and was prepared to explain how he worked. The problem was that he was required to address combinatorially massive problems with little evidence of

131

any real basis for his decisions. For instance:

> "What is the likelihood of this man, having this specific medical
> history, will suffer some illness as a result of his particular
> occupation or recreation and how would this affect his ability
> to do his particular job and for how long at what expected
> future salary over this period?"

Perhaps unkindly the impression was gained that some decisions were made on a rather arbitrary basis. However, since the circumstances of each case were so different it is unlikely that the premiums would be checked against each other. In effect the expert was acting as a "wise and unquestioned authority" (since someone had to set premiums) but, apart from general principles and a few well known cases (often involving the consequences of back problems) most of the decision making was rather "unscientific" and most cases were unique cases from an immense repertoire of possible combinations of health, occupational and other circumstances.

Permanent Health Insurance is a potentially very worthwhile area for the study of problem solving. Nevertheless, as a commercial prospect, the author found it was impossible to convince himself or the sponsors that a system could be implemented and the investigation was abandoned.

Additional Comments

The author's overall reaction to this task is that it is one of the more difficult that he has investigated. Apart from the skill in assessing the health and environmental risks etc, there was a real requirement to reason

from first principles on the impact of environmental factors on any health risk condition and, then, to assess the potential risk of disability on employability. Unlike many other tasks investigated, the possible combinations of factors are so great that there is limited opportunity to call upon the experience of previous cases.

Although the task was difficult (and perhaps because of this factor), the limited opportunity to investigate it did not leave the impression that it was particularly well executed. Or, more correctly, it was not obvious that it was well executed. Although, as stated above, there was a need to reason from first principles, the expert did not have any real theory upon which to base his reasoning.

Instead, there seemed to be a universal reliance on "ad hoc reasoning", perhaps using rough rules of combination, for every case that could not be directly or indirectly related to a previous case. In short there seemed to be considerable "finger in the air" guesswork highly protected by a web of professional mystique. The anomalies resulting from this guesswork were to be found in those examples where nearly similar situations resulted in very different quotations. The expert freely admitted that such situations were not uncommon.

This is the only task that the author has investigated where he would have liked to recommend a test of the expert's abilities. Thus he would have liked to present a test involving some very similar pairs of risks to determine the accuracy (or, more strictly, reliability) of the expert's judgement. Alternatively, a retrospective scan of previously assessed cases might have been used to the same effect.

Had an expert system been implemented there would have been an absolute need to originate a "theory of permanent health insurance" to embed in it. There is certainly an inadequate set of examples to just use a recognition system. The development of such a theory, even if not used to build an expert system, would seem to be very much needed in this area of insurance.

3.10   LIFE UNDERWRITING

Sponsoring Organisation

The implementation was sponsored by a major UK life insurance company.

Introduction

Life insurance is increasingly popular particularly when associated with house purchase. The sponsoring organisation were faced with an increase in business and difficulties in recruiting skilled staff. Also as a result of competition from other agencies there was also a requirement to undertake a proportion of the underwriting of new business away from the company headquarters so as to offer immediate quotations and cover at the company's branches. These factors motivated the successful implementation of the life underwriting expert system.

Author's Role

The author undertook the initial stages of knowledge acquisition and scoping for this major multi-user implementation. However, although he managed the implementation project he was not responsible for the detailed implementation of the resultant system.   Therefore the following represents his investigation of how human experts undertake this task rather than the specifics of the commercial implementation.

Outline Task Description

The task, which is the province of middle-ranked staff in a life underwriting department, is to analyze a life insurance proposal, initiate further medical and other enquiries where necessary and prepare a premium quotation.

Form of Investigation

The core knowledge acquisition was undertaken with a senior life underwriter who was both experienced with the current methods employed and who was also fully briefed by his superiors as to the commercial requirements of the eventual expert system. Nevertheless, as described below in the commercial implementation section, there remained for some time a mismatch between the desire the managerial desire to obtain the maximum sensible transaction rate and the underwriter's desire to make the best possible evaluations of the lives insured.

Later knowledge acquisition, which was not undertaken by the author, was considerably time consuming and detailed and was largely conducted on a segmented basis to populate the previously well-identified areas of the knowledge base. Other experts were also used at this later stage of knowledge acquisition.

The later implementation stages were very much preoccupied with integration and performance issues and residual knowledge acquisition was entirely routinised and undertaken by junior staff of the sponsor.

Findings on Nature of Problem Solving

The task commences with a life insurance application form submitted by the applicant either directly or via an insurance broker. The expert scrutinises the various fields on the form (eg health, occupation etc) looking for any entries that will either necessitate referral for further enquiries (eg a medical check of the applicant) or factors that will preclude insurance being given (eg recent cancer etc).

If this scrutiny does not indicate refusal or referral the expert will either quote normal rates for the applicant's age and sex or will quote a rate based on normal rates with a percentage "excess" loaded to reflect a perceived additional risk.

The existing manual task is supported by a mainframe computer system where the core details of the applicant are kept on a database also used to record the expert's quotation or refusal. This computer system is fully integrated with the eventual insurance premium collection system for applicants eventually accepted for insurance.

Problem Structuring Props (Tools and Aids)

The life insurance application form is the primary document used to structure the problem solving. The computer system containing some of the details is mainly used just as a recording medium for the final decisions. The form was routinely annotated as part of the problem solving. Further, in explaining their activity experts would constantly use the application form, pointing to the various reply fields and indicating "hidden line"

interrelationships between them.

Generic Data

Physical generic data consisted of "life tables" indicating the actuarially appropriate premium rates for normal risk applicants of both sexes and various ages. Subsidiary tables were also used to determine the correct percentage excesses for particular risks and details of those risks which must be met with a refusal to quote.

Risk Applicant Stereotypes

Another form of rather unique "generic data" was used to support this task. It was not physical data like that found above or in many of the other tasks but consisted of a shared set of "risk applicant stereotypes".

Underwriter training is largely by professional examinations and "apprenticeship" to more experienced underwriters. During this apprenticeship and during their ongoing career, underwriters come to share a set of stereotypes such as "thin teacher type" or "over 40 occupational drink risk type". Underwriters discuss applicants with their colleagues in descriptive stereotypes or variants of these. The stereotype descriptions are also occasionally used as annotation on application forms to help in the classification process.

The first of these stereotypes, the "thin teacher type" refers to an endomorphic teacher (or person in a similar social or occupational category)

138

who is likely to suffer from depression and therefore has an increased risk of suicide. The second refers to an applicant with a fairly high alcohol intake who is occupationally very exposed (eg barman or salesman) to alcohol as part of his job and, as a result, has a higher risk of cirrhosis of the liver.

Although the set of stereotypes (including many male homosexual stereotypes, such as hairdressers and "performers", suggested to be at risk from AIDS) is not written down it forms a powerful set of data (perhaps which can be considered as a problem solving prop) upon which certain categories of applicant are regularly classified. Such classifications sometimes result in a direct refusal or loading but, more often, forms a hypothesis which is later confirmed or rejected on the basis of medical or interview reports.

Specific Data

The application form, apart from its problem structuring role, also forms the primary specific data on the case. Secondary applicant-specific data is provided by medical and interview reports where these have been sought as a result of concerns raised on the basis of information on the application form.

Knowledge Sources

Three separate fields of knowledge are used in making the quotation decision. First, and supported by tables, actuarial knowledge is used to

assign the normal premium for applicant age and sex. Second, knowledge concerned with the interpretation of medical history declared on the application form, or of reports from medical examinations, is used to determine medical risks. The third field of knowledge concerns risks that might ensue from the applicants environment or activity. Thus occupational risks, risks due to sports or hobbies and risks due to residence in a foreign country are covered by this knowledge.

Open/Closed Problem

Life underwriting can be considered as an "artificially closed" problem. Thus, although all potentially insurable life risks are presented to the expert, he will only pursue those which can be categorised as "normal" or "near normal". Thus insurance will be refused unless an applicant has no apparent risk or possesses only a mild risk (or risk combination) which can be easily fitted into one of a pre-existing set of categories. Thus the potential "openness" of the problem is removed to leave just a residual closed set of risks each with its associated premium excess.

This is in contrast to Permanent Health Insurance where, although still strictly closed, the problem is much more like an open problem and includes many of the higher risk categories and combinations artificially excised from Life Insurance Underwriting.

Proactive/Reactive Problem

The problem solving is solely reactive to the presented details on the

application form.

Problem Solving Architecture

In Permanent Health Insurance there was the possibility of any one of an immense number of "diagnoses" resulting from the combinations of the multiple primary searches.   Life Underwriting has an apparently similar starting point to Permanent Health Insurance.   Also, and usefully for comparison, the exemplar tasks, Life Underwriting and Permanent Health Insurance cover very similar domains.   Both assess health risks and their potential future consequences to an individual.   In Life Underwriting the risk is assessed in respect of death.   In Permanent Health Insurance the risk is primarily assessed in respect of inability to work and obtain income.

However Life Underwriting and Permanent Health Insurance differ widely. In Life Underwriting there is artificial (ie institutionalised) constraint both at the primary search conclusions and at the "grand" final conclusion. These constraints serve to very effectively limit the effect of combinatorial explosion and make the task manageable and reproducible on a routine basis.

In Life Underwriting six of the search sub-domains are common to those in Permanent Health Insurance.   (The seventh sub-domain, Occupational Need, is not relevant since death is 100% occupationally disabling.)

Nevertheless, although the sub-domains are in common the search of each results in very simple "diagnoses" which are combined largely on the basis of maximum sub-domain risk to form the overall accept, adversely rate or reject conclusion.   A few combinatorial conclusions crossing more than one

sub-domain are allowed for on an individual basis. Thus an occupation of miner coupled with age over forty-five would trigger a special loading due to the extra risk of death from pneumoconiosis in older miners.

Life Underwriting is illustrated on the next page.

LIFE UNDERWRITING

Diagrammatic Representation of Problem Solving Components

Following the now common component of Global Data Acquisition, there is independent sub-domain search in the fields of Health, Job and other risks. To a large extent the conclusions from these searches are constrained to simple bands. Combination of the risks is also constrained and, rather than face a massive combinatorial problem as with Permanent Health Insurance, the Life underwriter merely insures the life on the basis of the worst loading from any of the sub-domains. If any unwelcome combinations occur then, unless they are common simple cases such as those discussed above, the proposal is usually rejected.

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusions	Secondary Conclusions

Multiple Domains with no Reference Linking

Constrained Primary Conclusions

Constrained Secondary Conclusion

Commercial Implementation Details

A system was implemented on an IBM 3090 mainframe under IMS first using IBM's ESE (Expert Systems Environment) and was later ported, to improve

multi-user performance, to Aion Corporation's ADS. It is successfully in use supporting over a hundred simultaneous users. The following factors were important for the commercial success of the implementation.

In terms of the actual expert system itself, it took some time to appreciate that senior management were aware of a tradeoff between the effort to underwrite a life proposal (in terms of the grade of underwriting staff involved and time taken) and the accuracy of the proposal evaluation. In the main it was the intention to underwrite most lives on standard rates (or commonly used special rates to allow for common medical conditions or risky hobbies/occupations) with the minimum effort but to ensure that special "hidden" risks, such as AIDS or alcohol related illnesses, were properly identified and the applicants referred for further medical examination.

This understanding was not properly appreciated by the underwriters who considered that they had a duty to make the best possible evaluation of every life even if this slowed down their work. In practice the management view makes sense since the less serious risks even out over the large number of lives underwritten. As a result the expert system had to offer a very rapid transaction rate for the majority of lives with referral for medical examination of the minority of suspicious cases.

The system had to be provided on the mainframe and had to be seemlessly integrated to the database and other software used to manage the life business. For instance, once a rate had been established, this had to be passed to the system that generated the offer letter to the applicant and to the system that would eventually collect premiums over the working life of the applicant. There was thus considerable implementation work required to integrate the system with existing mainframe applications and databases.

Also, since the mainframe application supported a conventional "green screen" dumb terminal it was essential to provide a compatible interface to the expert system. Thus users were presented with their usual low quality form of interaction with the mainframe and the appearance of the expert system did not conflict with the appearance with the rest of the underwriting system.

The expert system implementation was reasonably straightforward in terms of complexity but took rather longer than planned due to the comparatively low productivity of the mainframe tool used.

A serious problem arose as a result of the use of IBM's ESE mainframe expert system shell. This worked very well for single users but the lack of re-entrant code, both to support the inference engine and the knowledge base, resulted in the abandonment of ESE in favour of Aions's ADS. There 'is' not much experience of mainframe transaction processing expert systems relying on the mainframe engine.

Most expert systems applications involving the mainframe rely on an attached PC or UNIX workstation supporting each user as an intelligent terminal or, as in the well known American Express credit card application, rely on the mainframe accessing a "dark screen" attached workstation to deal with the minority of transactions requiring access to the expert system. The life underwriting system requires use of the expert system to complete every transaction and, with in excess of fifty concurrent users, special considerations in respect of multi-user support had to be seriously addressed.

146

Additional Comments

As identified above, one of the prime characteristic of this task is that there is a tendency towards expediency at the expense of accuracy or completeness. Thus, many theoretically insurable applicants are refused cover solely because their risk situation would take too long to properly assess. Similarly, quotations are usually made in only three or four bands, ie normal, +10%, +20% etc. This fairly course categorisation is adequate where actual risks are averaged between a very high number of insured persons.

The coarseness of the categorisation and the refusal of all except normal or near-normal risks enables the company to deal with a massive number of applications using only moderately skilled staff in the underwriting department. The fact that the expert system follows this philosophy was a key to its success and it is necessary to identify this "rapid satisficing requirement" factor in similar applications.

13.11 SMALL BUSINESS GUIDANCE

Sponsoring Organisation

This investigation and resultant system was sponsored by a Government Agency wishing to offer a free public service of guidance to small businesses.

Introduction

Small businesses (defined as involving between five to one hundred persons) suffer a high failure rate and the intention was to point to weaknesses and opportunities, particularly those that could be resolved by training.

Author's Role

The author undertook the requirements analysis for the proposed system, contributed to early knowledge acquisition and led and made the major contribution to the team undertaking the system design. He was not personally responsible for the later detailed knowledge acquisition or the full system implementation.

Outline Task Description

The task is to make a report on a small business. The task expert is usually an experienced accountant with special knowledge of the problems

of small business. The task involves diagnostic, advisory and referral skills. Thus the expert may discover or quantify a business problem, offer immediate remedial advice or, if the problem is specialised and outside his experience, make a referral to a further specialist.

Form of Investigation

Four principal and several subsidiary experts were employed in this investigation. This was because the range of experts' style and knowledge varied considerably more greatly than in any of the other tasks investigated. There is no accepted training or background for the task of small business guidance and, although most of the practitioners are accountants, they vary widely in the way they undertake the task.

To be specific, one expert could be best described as a "theoretician" with a well-developed theory of how small businesses should function. Another, whom we refer to as the "counselling expert", was particularly able to conduct a sympathetic and effective diagnostic dialogue with the owner-managers of small businesses. ("Owner-manager" is the term used to cover the principals of small businesses whether or not they are sole traders, partnerships or limited companies.) Such entrepreneurs can often be quite difficult to deal with and the expert was able to maintain a dialogue that held the expert's attention and respect.

The final two principal experts and the subsidiary experts were more generalist and had between them a wealth of experience on the sorts of problems and unexploited opportunities faced by many small businesses.

The method of knowledge acquisition was adapted to the different areas of expertise represented by the four main and the subsidiary experts. The "theoretician" was interviewed in a rather abstract way with only passing references to individual cases to illustrate points in the theory of small businesses.

The two generalist principal experts and the subsidiary experts were interviewed in a manner very much like that for the other tasks investigated. There was thus considerable use of complete and partial case details and, in the later stages of knowledge acquisition, considerably exhaustive knowledge acquisition on particular business problems and their resolution.

The "counselling expertise" principal expert presented a particular challenge. It was obvious that an expert system could not offer the degree of personal understanding that the expert was able to provide. Thus knowledge acquisition had two motives. First, the expert's own method of conducting a dialogue was studied and, second, the expert was asked to advise on the special needs of small businessmen in the interview situation.

The results from the counselling expert were interesting and challenging from the implementation aspect. Most of the skill involved listening to the owner-manager with only small prompts to acquire more detail or direct the flow of the owner-manager's flow of conversation. If properly handled, owner-managers need little encouragement to tell the history of their business from its foundation to its present state. As a "listening" expert system is currently not feasible, it was necessary to devise substitute functionality in the form of a more traditional expert systems dialogue that would hold the owner-manager's attention for what was realised would be

a rather long diagnostic and analytical session. The results, involving enhanced feedback and "route maps" is described below. However, the counselling expert was able to advise on the needs of businessmen in respect of such feedback.

First, it seems that owner-managers are rather insecure and often quite impatient. As a result, the counselling expert advised that she offered considerable positive feedback to assure the expert on those aspects where their business was in good order. Second, she considered it was essential for her to offer this feedback during the course of the dialogue rather than in the final summing-up or report. Third, and similarly, she found it essential to offer any intermediate conclusions during the course of the interview rather than wait to its conclusions. Fourth, it was particularly useful for her to explain the progress of the interview so that the owner-manager realised what aspects of his business had been covered and what was next, and to come, on the interview agenda.

The counselling expert was very clear on the above needs of the owner-manager and they were confirmed with the other experts.

Findings on Nature of Problem Solving

Expert guidance was offered as a result of a two-stage process. First, a general "top-level" scan of the business covering all of the business areas is conducted. During this scan, and usually the first part of it, the main parameters concerning the nature of the business, its size in terms of personnel and turnover etc are collected.

This top-level scan serves to highlight any perceived existing problems or unusual features and acts as input to order and focus the detailed second stage. During this stage each of the business areas (eg Finance, Premises and Equipment etc) is visited in turn and subjected to a more-or-less isolated separate analysis.

Finally, the expert reports on particular problems perceived in each of the business areas and offers general observations on the overall state of the business.

Problem Structuring Props (Tools and Aids)

As already described the experts in this field are particularly responsive to the facts as offered by the owner-manager. Thus structuring, apart from the two stages just outlined, is largely ad hoc and there are no real props or aids to structure the problem solving. However, all experts take notes during the course of their interviews with owner-managers and, to some extent, use these notes as aids to structure the subsequent parts of the session.

In the examination of the financial area of the business the accounts are sometimes used as a starting point to structure the examination. Nevertheless, these usually fall into the background as problem areas begin to emerge. For instance, although the accounts can be used as a quite neutral "conversation opener", the entries on provision of pensions for the owner-manager can lead to the very sensitive area of training for succession when the owner-manager retires or dies. This area is often one of the most important to consider since the independent nature of many

small businessmen makes it very difficult for them to even consider who will take over their "baby" when they retire. Neglect of proper provision for succession is one of the main reasons why small businesses fail after the unexpected departure or death of the owner-manager.

Generic Data

The business guidance experts are primarily looking for potentially dangerous abnormalities or unrecognised or unexploited opportunities in the business. To do this they must have a good understanding of what is currently "normal" for a business of the type and size under discussion. Some of this comes from their experience of examining many types of small business. The remainder, which is the generic data upon which they rely, comes from specialist business sector reports and from more generalist financial newspapers and journals.

From the business sector material they can determine what is acceptable in terms of say, liquidity, for a electrical wholesaler (where stocks may not be easily disposed) or a service organisation (where there are little stocks). From the generalist material (eg the Financial Times etc) they keep an understanding of the short-term effects likely to affect all businesses. Both these sources of generic data were claimed to be of considerable use to the experts in their guidance work.

Specific Data

Specific data on the small business under examination is primarily elicited

from the owner-manager. Unlike larger public companies, small business often produce little more than statutory accounts, simple sales literature and price lists.

Knowledge Sources

Seven knowledge sources covering separate fields of business knowledge can be identified. The principal knowledge source covers financial aspects of a business with personnel, marketing, production, distribution, suppliers and premises/equipment representing the other areas of knowledge. These knowledge sources and their interrelationships are covered in the discussion on problem solving architecture.

Open/Closed Problem

Business guidance is an open problem with an infinite number of potential business problems (or opportunities) and corresponding solutions. Thus in the current (and any near future) implementation it is not possible to entirely replicate the expertise of the human practitioner.

Nevertheless a high proportion of cases fall into established problem or opportunity categories and it was therefore possible to represent these in the implemented system. Further, limited reference to a generalist or specialist advisor was possible for abnormalities which were recognised but outside the advisory scope of the system.

An Example of Problem Referral

The above described system supported referral is entirely analogous to the current human situation where the business advisor realises that there is a problem in, say, the firms transport and distribution system. However the generalist advisor is unable to fully quantify the problem and offer remedial advice. In these cases referral to a distribution expert is made.

Referral of partly described problems within networks of interrelated expert systems (or human experts) is an important topic for future research. Such referrals are a key element of internal business mechanics and are only just starting to be addressed by the expert systems workers.

Proactive/Reactive Problem

The problem solving is largely proactive to the adviser's objective to discover problems and opportunities in the business and answer these with solutions to the betterment of the business. Nevertheless, as the dialogue continues, the system can be viewed as reactive to the problems and unanswered opportunities elicited from the owner-manager.

Problem Solving Architecture

Small Business Guidance can be illustrated as on the next page:

SMALL BUSINESS GUIDANCE
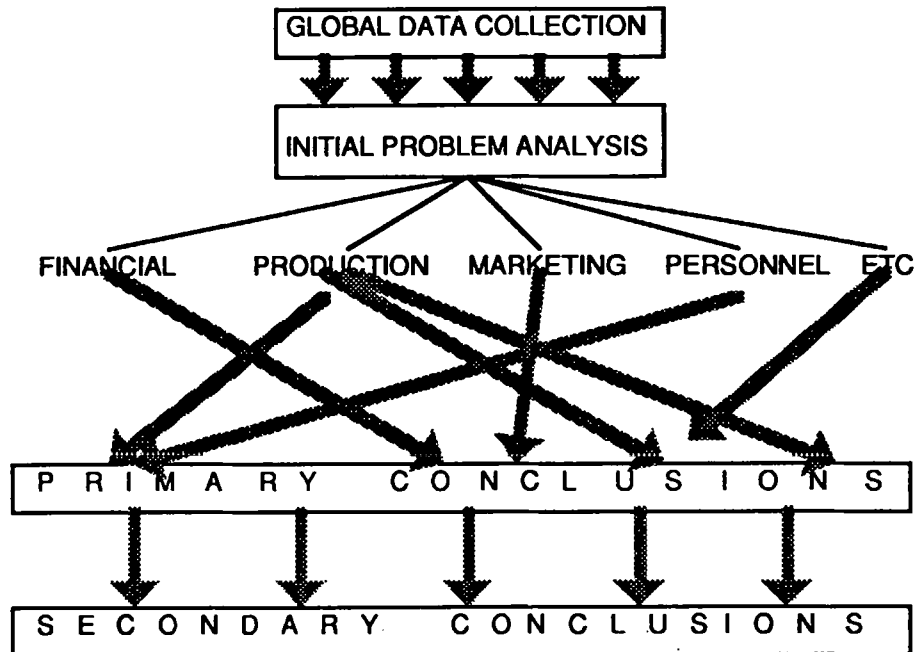
As described above Small Business Guidance is characterised by an initial problem analysis search. This initial search highlights the immediately obvious problem areas of the company and selects and prioritises the order of those sub-domains for subsequent detailed search and analysis. In the simplest case, for example, the areas for attention will be identified directly from the owner-manager seeking advice. He will usually know if his main problems are in, say, Marketing and Distribution, and so these sub-domain topics will be searched first.

In Audit Work Planning and Business Performance Assessment the sub-domains were sufficiently isolated for genuine separate serial searches to be undertaken. However, in Small Business Guidance, Job Shop Scheduling, Eurobond Trading etc it is not usually possible to complete the search of

157

any one sub-domain without referencing, to a greater or lesser degree, the information or partial conclusions proper to one or more of the other sub-domains. We refer to this inter-sub-domain access as "reference linking".

This "reference linking" does not indicate that the sub-domains are fictitious and that a single continuous domain exists. Quite definitely the experts consider that they are "now investigating Marketing" etc and the lateral references to allied sub-domains are merely to support the search of the current "focus of attention" sub-domain (ie Marketing etc). Reference linking will be a recurrent topic in the descriptions of the tasks which follow and, for this reason, it is discussed separately in the next part of this report.

In the discussion of Audit Work Planning and Business Performance Assessment we showed that, following essentially separate searches of allied sub-domains of a single task, elements of the individual sub-domain conclusions were combined to form secondary conclusions. In Audit Work Planning such secondary conclusions were fragmentary and of minor importance to the primary sub-domain conclusions. In Business Performance Assessment the secondary conclusions were a major outcome of the primary conclusions and were of paramount importance.

Similarly, in Small Business Guidance, following the search of each of the business sub-domains in the order chosen as a result of the initial problem analysis, the primary conclusions will be combined to make some general conclusions on the state of the business. However, unlike Audit Work Planning and Business Performance Assessment, both primary and secondary conclusions are of equal value in Small Business Guidance and neither can be considered the dominant product of the guidance session. Therefore, a

158

small business advisor may make some very pointed observations on the need for training within the firm and combine these with some general conclusions, accumulated from analysis of all the business functions, on the overall management of the firm.

Sub-Domain Reference Linking

Sub-domain reference linking may be illustrated diagrammatically in the way shown overleaf. To provide a better basis for understanding of this and the following tasks, this representation and the following explanation are in far more detail than the more usual diagrammatic representation of problem solving architectures and their constituent components.

SMALL BUSINESS GUIDANCE REFERENCE LINKING
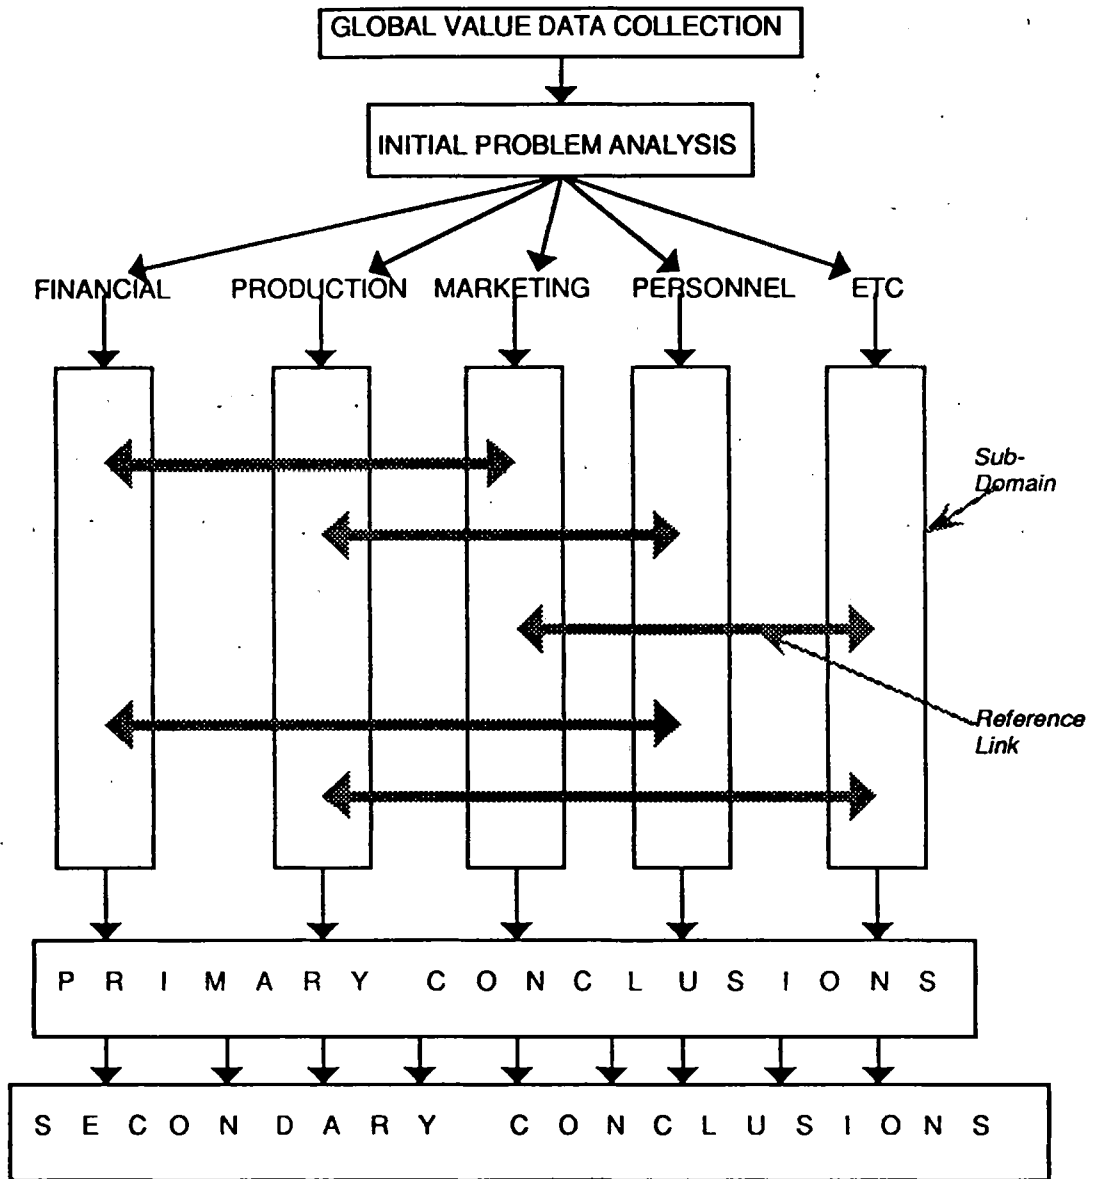
Most of the inter-access is to and from the financial sub-domain. Thus, for instance, part of the expert's consideration of the firm's personnel situation involves an assessment of the residual adequacy of company skills if a key employee (who may be the proprietor or a director) should suddenly leave the firm. This check is essential if the future success of the firm is not to be critically dependent on one individual.

If a potential weakness is discovered, the expert will move toward a recommendation to recruit backup skills. Since this will involve an element of growth (ie increase in staffing) the expert interrupts his analysis of personnel matters to consider the firm's financial situation to see whether growth is anticipated. This is usually the case in a healthy firm and, on the basis of a good indication on financial matters, the expert will proceed to make a recruitment recommendation before going on to finish considering the other personnel matters.

If, for some reason, the firms financial situation shows that an increase in staff cannot be tolerated, the expert returns to complete personnel considerations looking especially for low-cost options to cover the identified skill vulnerability. This analysis results in options for deferred recruitment, skills training of existing staff or, in the worst case, a warning on skills vulnerability in the concluding report.

The expert's access from the personnel sub-domain to the financial sub-domain can take one of two forms. If, as a result of the initial problem analysis, the expert has chosen to consider the financial sub-domain first, then the later consideration of the personnel sub-domain can reference (ie "draw upon") the established primary and inferred information in the "completed" financial sub-domain.

If, on the other hand, the expert's initial problem analysis results in personnel being given first priority in the order of sub-domain considerations, there will be no pre-existing inferred information available from the financial sub-domain.

In this case, when the expert breaks his personnel considerations to

reference some financial information, just those inferences and computations necessary to "answer" the current personnel considerations sub-domain enquiries will be made by expert in the financial sub-domain. However, the results of such work will be recorded or remembered by the expert until a full consideration of the financial sub-domain is made. This avoids repeated work and, in some cases, repeated access to the businessman consulting the expert.

Expanding on the example given above, if personnel problems are shown to be paramount as a result of the initial problem analysis, then recourse to the financial sub-domain for "growth permission" will be to an as yet unexplored sub-domain. In this case the need for early financial information does not prompt the expert to make a complete analysis of the financial sub-domain.

The expert in such cases usually makes two "quick and dirty" checks on financial performance (current profit to turnover ratio and growth curve over last five years) and, on these two crude but powerful indicators, decides whether the firm can support the need for personnel increases. Since the financial checks themselves form an element of the more complete financial analysis, their results are stored to support that analysis if, as is likely, it is required.

Conversely, as is usually the case, if there is an analysis of the personnel situation following analysis of the financial situation, the personnel search can make recourse to the full considered conclusions on growth.

References between other sub-domains are possible. For instance, can a product identified as "under-marketed" be actually produced. This entails

a reference linking from "Marketing" to "Production". Such reference linkings, although often very important, are rather less common than references to the "Financial" sub-domain.

Diagrammatic Representation of Problem Solving Components

The diagram shows the presence of the new Initial Problem Analysis component plus link lines to show the minimal inter-sub-domain reference linking just described. This diagram is rather redundant in view of the more complete diagram above but is shown to indicate the progression from previous tasks and also because the more complete form of diagram was solely used on this occasion to explain the concept of reference linking.

Multiple Simple Search with Minimal Reference Linking

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Minimal Reference Linking

Limited Overarching Conclusions

Iterative Feedback
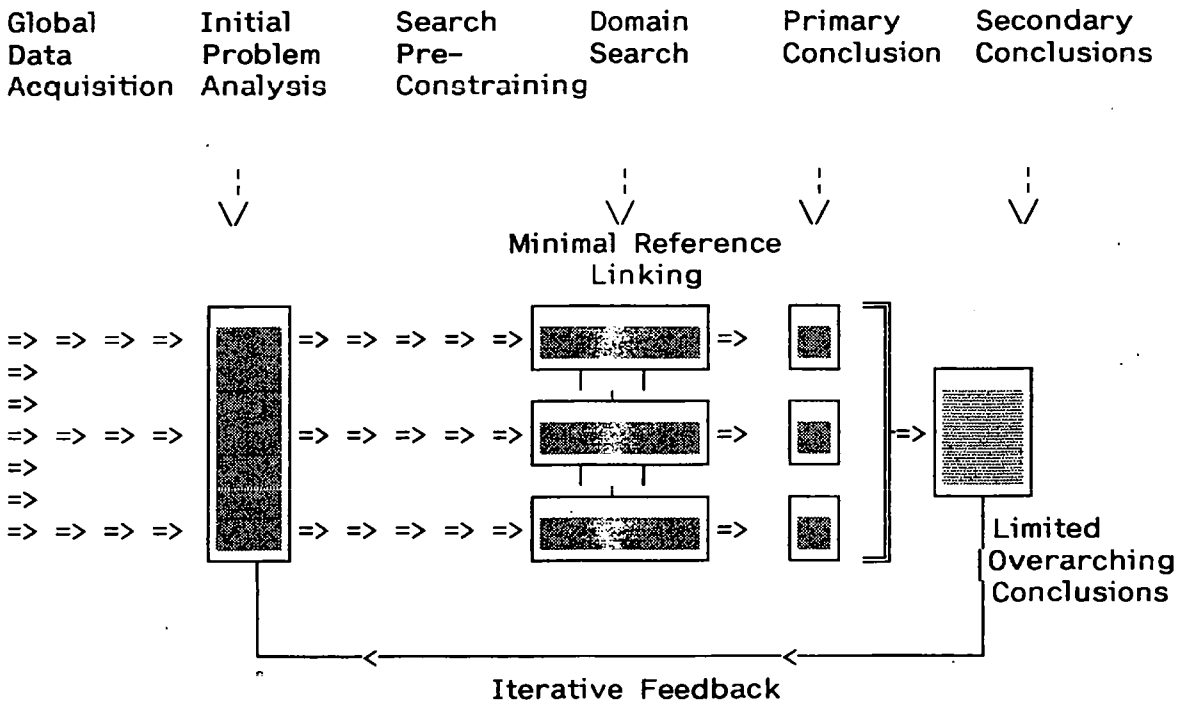
Commercial Implementation Details

This was a successfully implemented and operationally deployed system. Government funding to the sponsor was a key factor and it was possible to implement a system with a proper budget and with the help of enthusiastic sponsors.

The author was primarily involved in the pilot but the full system was just a more complete version with increased detailed knowledge on more aspects of business.

The system was implemented as a large system using the Crystal 3 expert system shell. Although the basic form of problem solving is based on the experts' knowledge, the implementation was partially on designed functionality rather than being a replication of the experts' skills. This is mainly because, as already explained, the human expertise involves considerable listening skills and this cannot be replicated in an expert system.

The system would not have been successful if we had not been alerted to a key characteristic of the intended user population. The sponsor and counselling expert very much understood the characteristics of the proprietors and directors of small businesses. Many of these were rather idiosyncratic people who worked on their own account because they did not fit in with the discipline and procedures of large corporations. It was claimed by the sponsors that they had a relatively short attention span to new technology and, if a system did not grip their interest within about ten minutes, there was little hope of getting them to adopt it in their business. In practice this proved to be good advice and the guidance system was

designed with this constraint in mind.

If business guidance is to be properly produced each of the business areas must be fully explored on a sequential basis, with minimal linkages (usually to the financial area), and a complete recommendation given to the small businessman. This process takes considerably longer than ten minutes and can require the businessman user to input a considerable amount of data. Therefore the system was designed to work on a "two pass" basis. An initial interactive dialogue, requiring straightforward answers, produced an interim diagnosis of the key area where guidance was needed. This feedback to the user was usually sufficient to capture his interest so that the more exhaustive analysis and diagnosis could be made.

It was also, incidentally, a very useful first stage since it allowed the first pass results to direct the order in which the business areas were searched for the second pass. Without the initial diagnosis the system might involve considerable user input and spend much time on financial matters when the real problem was in the area of providing adequate funding for the marketing of a new product. With a first pass diagnosis the search is reordered from the default order (starting with financial analysis) to deal with sales and marketing matters first.

The delivered expert system conformed very much to answering the "insecurity and impatience" needs of the expert as advised by the counselling expert. Gratuitous feedback was liberally offered for those cases where everything appeared in order and intermediate conclusions were displayed as text screens. A "route map" of the business sectors was also provided and completed sections were successively coloured in red to show that the topic had been concluded and that the user has "successfully"

completed that stage of the interaction.

Hardware considerations were also an issue with this system. Although the first pass diagnosis might be undertaken at the agency premises, the system would only be useful on a long-term basis if it could be used on the small businessman's own PC. Privacy was an important issue and, additionally, it was anticipated that the businessman was more likely to use the system in the evenings and at the weekends when he was not busy running his business.

It was thus a condition of the implementation that a run-time version of the system should be made available on floppy disks for use at the businessman's address.

Additional Comments

As already discussed the implementation is very different to the human expert situation where the business advisor is sensitive to clues arising from the conversation with the owner-manager. The implementation nevertheless is rather similar in style to the pedantic expert who uses a more structured form of interview or the interview with a "monosyllabic reply" owner-manager.

As a result the expert system implementation tends to be more protracted than the expert interview and, as discussed above, special care has to be taken to provide encouragement via feedback and route maps.

The current system could be ideally extended to be more sensitive to

different business sectors. Currently sector-specific knowledge is incorporated as part of a generalist knowledge base and this is both unwieldy and inefficient. An improvement would be to implement a wholly generalist system which could detect that the business was, for instance, a building contractor or retail shop and load an appropriate knowledge base to continue with a sector-specific consultation.

Similarly there is a requirement to deal separately with very small businesses (ie sole traders up to about four persons) where it is generally accepted that problems are very different from the larger small businesses where there is a recognised internal structure, responsibilities and means of decision making.

3.12   PENSIONS MANAGEMENT

Sponsoring Organisation

This investigation and the subsequent commercial implementation were sponsored by a major distribution company with regional offices and pensions staff throughout the UK.

Introduction

Occupational pension schemes are complex. Often one or more new schemes are brought into use during a scheme member's working lifetime. The interpretation of such schemes and, in particular, the benefits accruing to individual members is often difficult and time consuming. This task was investigated in order to provide expert systems assistance to pensions staff responsible for answering members' queries.

Author's Role

The author was responsible for the instigation of this expert systems project, initial knowledge acquisition and analysis and for the later stages of the management of the commercial implementation. He was not responsible for detailed knowledge acquisition or detailed final system design.

Outline Task Description

The task is undertaken by two grades of staff. Primarily the task is the responsibility of Pensions Officers. These "professional staff" vary in their knowledge of the relevant current and historical pensions schemes. Some Pensions Officers are Scheme Trustees with first-rate knowledge of the schemes. Others are rather less knowledgeable and effectively still trainees.

The second grade of staff to undertake the task are Personnel Officers who have far less acquaintance with the schemes but have to answer routine pensions questions as part of the personnel function of the organisation.

Although different grades are involved, and deal with questions of greater (Pensions Officers) or lesser (Personnel Officers) complexity, the task is much the same in both cases. Queries are initiated by scheme members usually at critical phases of their career. These are shortly after recruitment, on marriage, on divorce, just before or after retirement and, by members' dependents, on death.

The queries are either general (eg "How is my lump sum calculated") or, very much more often, specific (eg "Will my new wife be covered by the contributions which I paid for my divorced ex-wife and, if so, how much pension will she get if I die?"). The task of the expert is to give an authoritative answer to the member explaining the basis of any benefits together with a calculation where necessary. Where required the expert must give references to relevant scheme paragraphs or previous correspondence extending the coverage of schemes. (However the discussion on "Authoritative Text" later in this section should be noted in this connection.)

Form of Investigation

No unusual knowledge acquisition techniques were used and the normal structured stages described elsewhere proved very productive for a very extensive and detailed knowledge base. However, although the knowledge acquisition was straightforward, there were considerable difficulties in reconciling the knowledge and opinions of different experts. This issue is discussed under "Commercial Implementation" later in this section.

Findings on Nature of Problem Solving

As described below the structure of problem solving is not particularly complex. Further, the problem is usually constrained by the member to just a specific aspect of the pensions scheme and member's situation.

Nevertheless, a prime characteristic of problem solving in Pensions Management is that it relies on a very large knowledge base. Problem solving may not be complex but the knowledge available is very voluminous. Schemes and their associated documentation run to hundreds of paragraphs and such material is incorporated as part of the expert's knowledge. The special relationship between the scheme knowledge incorporated into the expert system and the actual text of the schemes is discussed below.

"Authoritative Text"

Pensions schemes are a set of regulations. As such they are the "laws" under which pensions are paid. Legal and quasi-legal expert systems, such

as that developed for Pensions Management, may run into particular problems on what, in the case of Pensions Management, became known as the issue of "Authoritative Text".

At the outset, the requirement for Pensions Management was for a conventional expert system offering guidance on pensions schemes queries. As the system became successful the requirement changed in favour of a system that actually delivered advice using the precise wording of the schemes. In other words, the expert system was required to deliver advice more-or-less at the level that could be supported in Court.

The precise outcome of this requirement is not strictly relevant. In fact, output text was produced which was compatible with the text of the schemes but was not guaranteed to be "authoritative".

The general issue is however relevant for "legal advice" expert systems. Where the requirement for "Authoritative Text" becomes paramount then the knowledge base of the system become an image of the underlying "law". For example knowledge base rules map directly to the legal rules. There is then a danger (and the suggestion was actually made for the Pensions Management expert system) that the knowledge base becomes the authoritative document in which the law (ie scheme) is embodied. We consider that the concept of a "runnable law" is interesting but presents many real problems both for the lawmaker and the knowledge engineer.

This issue will not be further discussed but is presented purely because it actually arose as a serious issue in the implementation of Pensions Management.

Problem Structuring Props (Tools and Aids)

The human expert bases his problem solution around the letter he will prepare in response to the member's query. The fairly rigid structure of such letters acts a "virtual" problem solving prop rather similarly to the quotation prepared for Commercial Business Underwriting.

Generic Data

The Pensions Schemes manuals and supporting documents, such as Pensions HQ Advisory Circulars formed the generic data supporting the Pensions Officer.

Specific Data

Members records, including financial contributions and elective options such as to make extra contributions, form the specific data upon which advice is given.

Knowledge Sources

Pensions knowledge covers several fields which can be easily considered as separate knowledge sources. Nevertheless the "discreteness" or separation of the fields is not so marked as in, say, Small Business Guidance to which this task is structurally related.

A particular and quite discrete knowledge source covers the granting of credits for service with previous companies or for service with the Armed Forces. This knowledge impinges on the problem solution in quite a different way to the other knowledge employed.

Similarly knowledge concerning the eligibility of "children" is particularly specialised. A "child" may be a dependent student in his or her late twenties but may not, for instance, be an infant born from a member's secondary relationship. Such knowledge, complex but totally separate from the main body of pensions knowledge, merely serves to define the number of children in the main benefit's evaluation.

Other knowledge sources, such as that covering Widow's Benefits, seem closer to the main body of the pensions knowledge and their separation into discrete knowledge sources seems more artificial.

Open/Closed Problem

Pensions Management can be considered a closed problem where the scheme should cover all possible eventualities. Nevertheless, in the case of the human Pensions Officer, there was always a trickle of novel cases which, however, were directly referred to the Scheme Trustees for arbitration. Such cases can therefore be considered outside the competence of the expert whose expertise covers a closed world of problems and solutions.

Proactive/Reactive Problem

The task is solely reactive to specific enquiries from scheme members.

174

There is no objective other than to answer these queries.

Problem Solving Architecture

Pensions Management has much in common with Small Business Guidance. In fact Pensions Management can be best regarded as a special simplification of the architecture of Small Business Guidance. Pensions Management has also something in common with Multinational Tax Planning and Engineering Preventative Maintenance in that the initial problem is constrained. In the specific case of Pensions Management the constraint comes (via the pensions officer or clerk) from the scheme member who, in most cases, just requires clarification on a particular aspect of his pension scheme entitlement.

For instance, on re-marriage, a divorced member will often enquire about widow's pension entitlement. Similarly, a member nearing retirement will question whether service with a previous company, or with the Armed Services, will contribute positive credits toward the amount of his or her pension entitlement.

In all cases the search is based on retrieved records of member's contributions and options selected in the past. However the search is constrained to that sub-domain relevant to the member's enquiry. However, whereas in Multinational Tax Planning it is possible to constrain the search to, say, American countries without reference to Asian company matters, this is not possible in Pensions Management.

Widow's pension entitlement will be primarily based on the member's prior options (ie to include widow's cover on marriage and, perhaps, to cancel it

on divorce) but will also depend to some extent on other circumstances. If the member will have "children" at the time of his death then the widow's entitlement may be reduced to provide some allowance for such "orphans".

Thus there is some reference linking from the "Widows" sub-domain to the "Children" sub-domain to determine the children's status. This may be simple if there are no children or they are all independent adults. However, a more complex search of the "Children" sub-domain may be necessary if there are adult children who are "Students" within the meaning of the Pension Scheme. This may result in a detailed dialogue with the member over such matters as "academic" versus "craft" training and determination of full or part-time student status. In some cases a more-or-less complete search of the "Children" sub-domain will be necessary to provide the required values to allow the "Widows" search to conclude.

As with Small Business Guidance, the inter-sub-domain references are not heavy and the expert can provide a competent response to the member's enquiry. The member is asked to confirm that the response adequately meets his or her needs and, if so, the response can be regarded as final. However it is quite common for the reference searches to raise issues which the member had forgotten or disagrees with. Matters concerning credit for membership of previous schemes often arise in this way. In such cases the search will be repeated but constrained to the new area of enquiry. Naturally conclusions from the first search will be retained and may be referenced from the sub-domain search covering the new focus of attention.

In some cases, often in the period from five to ten years before retirement, the member will want a complete assessment of his pension and other cover and, when this happens, no constraint is put on the search. In such cases

all sub-domains are searched in a fixed order with inter-sub-domain reference proceeding exactly as in Small Business Guidance.

Diagrammatic Representation of Problem Solving Components

The diagram components have all been previously encountered but it is noteworthy that Pensions Management was the only task considered that involved an example of all of the pre-, post- and actual search components. The secondary conclusions are constrained since the majority of the requirement on the Pensions Officer is to give detailed advice in the form of primary conclusions following consideration of the individual sub-domains of the pension problem.

Pre-Constrained Search with Minimal Reference Linking to Other Search Areas

Global Data Acquisition   Initial Problem Analysis   Search Pre-Constraining (Domain Driven)   Domain Search   Primary Conclusion   Secondary Conclusions

Minimal Reference Linking

Constrained Overarching Secondary Conclusions
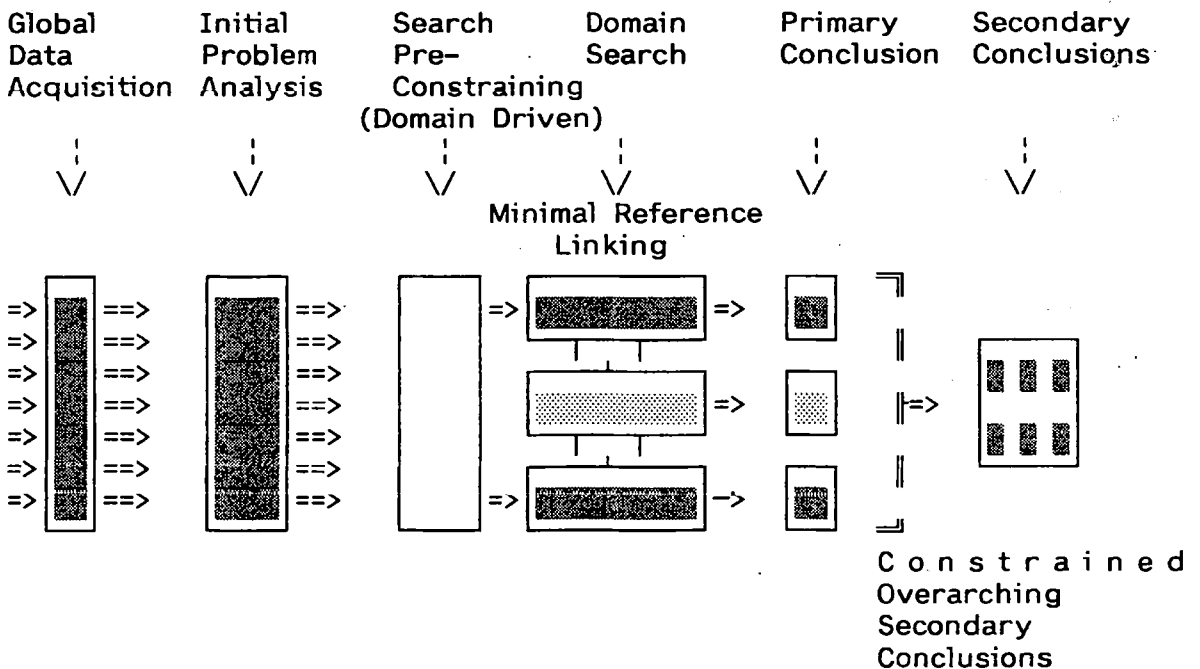
Commercial Implementation Details

The resultant expert system was implemented, not by the author, using the Crystal 3 expert system shell.

This project resulted from a competitive tender which was won largely on the ability to demonstrate the form of the user interface that the eventual system would support. The user interface issue was key to acceptability by the pensions officers who were to form the user population. A simple screen painter was used to demonstrate a chain of screens simulating the form of the interaction.

The project nearly failed due to disagreement between the experts. Each of the company's regions had its own chief pensions officer who was responsible for interpreting the rules of the pensions scheme. For instance, is an adult mature student following a part-time course a "child" when it comes to considering entitlement for a share of his fathers occupational pension death benefit? To what extent does service in various parts of the armed services count toward years of service under the scheme rules? Strangely, the answers to such questions differed across the country and beneficiaries received different treatment under the rules of the same scheme.

However, an expert system can only give one form of advice and the disagreement between the key expert who was the subject of knowledge acquisition (and on whose knowledge the system was based) and his colleagues caused time-consuming and expensive delays in completion of the system. There was also considerable criticism of the expert system by the experts who disagreed with its advice.

The issue was finally resolved by the sponsor agreeing that a single expert would act as mediator between his colleagues and that he would seek consensus and that the system would reflect this consensus knowledge. Although this seems a very commonsense arrangement it was actually very difficult to achieve and there were grumblings of disagreement between the experts right up to finalisation of the project.

Additional Comments

This investigation and implementation were straightforward and a clear understanding of the problem solving structure was ascertained. However as a part of "User Knowledge Acquisition", this task did illustrate the different needs of the different "Pensions Officer" and "Personnel Officer" user populations.

A requirement for different user interfaces to support the two populations was agreed for a later version of the system but the author was not involved in this work. The requirement for "Authoritative Text" was also associated with the needs of the more skilled "Pensions Officer" user population and has already been fully discussed above.

3.13 JOB SHOP SCHEDULING

Sponsoring Organisation

The investigation was conducted at the invitation of a UK based domestic electronics product manufacturer.

Introduction

The job shop consisted of a number of machines each capable of performing different manufacturing operations (eg cutting, drilling, assembly, testing) to assemble a range of manufactured goods of various sizes and complexity and in various quantities.

The job shop was order-driven and it was the task of the scheduler to maximise the number of jobs passing through the shop and to ensure that each met its due date etc. Increase in job shop utilisation was required not only to increase profit but to enable the plant to cope with expected growth in demand without the need for extending the job shop manufacturing facility.

Author's Role

The author solely conducted this investigation.

Outline Task Description

The scheduler and his colleagues receive regular details of incoming orders from the manufacturer's sales department. Each order has details of the type of product (which can be from a large range), quantity and due date. Other details may also appear on the order and, particularly, its origin is noted. Orders from key accounts are given priority in the case of any contention for manufacturing resources. Orders designated for stock maintenance are usually given the lowest priority.

The scheduler prepares a rolling schedule, extending for between one and two months into the future, as a Gantt Chart.

A secondary task of the scheduler is to make reactive changes to an established schedule so that rush orders and machine breakdowns can be accommodated in the schedule.

(For clarification it should be noted that the production plans (ie sequences of operations) for any particular type of product are fixed and that the schedule for any order is primarily a time and machine sequencing of a particular pre-existing plan. Of course Gantt charts are used for planning problems but, in Job Shop Scheduling, only scheduling is involved.)

Form of Investigation

The nature of the manufacturing was specialised, with rather unusual machines, and the author spent some time studying documentation and seeking advice from other known experts in the field before visiting the job

shop.

Initial knowledge acquisition consisted of a detailed tour of the job shop, during which notes were taken, followed by two extended knowledge acquisition sessions. The investigation has not yet resulted in an implementation and there is still further requirement to investigate the scheduler's problem solving. Nevertheless, sufficient knowledge acquisition was undertaken to complete the structural analysis of the scheduler's problem solving as outlined below.

Findings on Nature of Problem Solving

Scheduling is acknowledged to be a combinatorially difficult task. Particular difficulties arise in the case of near capacity schedules where the scheduler must arbitrate between a number of constraints. For instance, if it is unlikely that all orders can meet their due dates then which orders can be allowed to slip? Must machine maintenance always take precedence over completion of an urgent rush order? Is it worth elaborate re-tooling of an existing machine just to enable it to cope with an order that cannot be accommodated anywhere else on the shop floor?

All these issues, and several others, were encountered in the Job Shop Scheduling investigation. Most of the resulting knowledge acquisition (and that anticipated for any further investigation of this problem) concentrated on the way in which the scheduler copes with a "combinatorially impossible" problem.

From our investigation we confirm that human schedulers do not produce

optimum schedules. The resolution of the many conflicting constraints (which we describe below in the form of "heavy inter-sub-domain referencing") is beyond the capacity of human schedulers. As a result they have developed procedures to develop "good enough" (ie satisficed but occasionally considerably non-optimum) schedules on a routine basis.

Problem Structuring Props (Tools and Aids)

The primary scheduling workspace is the Gantt Chart which, in the Job Shop Scheduling task, was a wall-mounted large whiteboard with permanent markings for the days of the month etc. The schedule was assembled and modified on this chart using felt tip markers. Subsidiary paper forms and charts, showing individual machine capacity plans etc, were used by the scheduler on a table in front of the wall mounted Gantt Chart.

The Gantt Chart is a good example of a problem solving prop and (although actually a whiteboard) also very clearly illustrates human co-operative problem solving using a "blackboard architecture". The scheduling is actually distributed between several experts who have responsibilities for different parts of the job shop schedule. (This issue is separate from the "hierarchy of schedules" issue discussed under "Additional Comments" at the end of this section. The separate schedulers discussed here all contribute to the main "top level" schedule. This "top level" schedule provides the target dates and quantities for separately undertaken schedules further down the hierarchy.)

Such schedulers, for instance, are aware of the characteristics of different types of machines and can comment on how constraints may be relaxed in

184

an impasse situation. The Gantt Chart acts as a common workspace for the evolving schedule and the scheduling experts are each able to see how various parts of the schedule impact on their own part of the job shop. As a result they can solve problems, say, concerning "bottlenecks" in their own area or, if these cannot be solved, they can alert the other schedulers that some relaxation must be made in their areas to facilitate solution of the local problem.

Generic Data

Data on machine capacities, in terms of sizes of components which can be processed and throughput per hour, forms the primary "capacity" generic data upon which the expert bases his problem solving. Subsidiary generic data is in the form of machine maintenance schedules, available workforce for job changeover re-tooling of machines etc.

Specific Data

Order data is the primary specific data. Also required is data on work in progress or already scheduled for the current period. Thus a rolling schedule inherits some work and commitment from the previous scheduling period and, in turn, makes some commitment on the future scheduling period.

Knowledge Sources

Specific knowledge sources can be clearly identified for this task. (This

division is over and above any division anticipated to replicate the "whiteboard" contribution of separate scheduler's knowledge as described under "Problem Structuring Props" above.)

A key knowledge source concerns the scheduling of specific orders and the relaxation of constraints in respect of those orders. Another primary source of knowledge concerns scheduling of resources including both machines and staff supporting them. To a large extent, problem solving consists of an arbitration between these two knowledge sources and the eventual schedule is the best satisfaction of the orders within the available capacity.

Other knowledge sources deal with the separate issues of machine maintenance (although this can sometimes be considered as a mandatory "no product" task for scheduling), transport and storage of partially assembled products between machines and with the arbitration between different sets of constraints.

Open/Closed Problem

The problem is closed but the potential number of solutions (ie schedules) is combinatorially explosive.

Proactive/Reactive Problem

The problem is solely reactive to the need to produce schedules to satisfy orders.

Problem Solving Architecture

In Small Business Guidance (Multiple Simple Search with Minimal Reference Linking) just considered, the references to sub-domains outside the current focus of attention were relatively infrequent. For instance, it would be rather unusual for references to be made more than twice from any one sub-domain to another during the solution to a particular problem.

Job Shop Scheduling also has a series of sub-domains which can be considered as separate. However reference linking between them is very considerable and, to some extent, the expert can be viewed as considering a number of disparate sub-domains on a "co-operative" basis. That is, there are many requirements for information to be passed between sub-domains as each is considered and thus the sub-domains may be thought of as co-operating.

In Job Shop Scheduling (and Eurobond Trading which follows) there is thus the prospect of very heavy reference between sub-domains. In fact the referencing is potentially so heavy that human problem solvers fail to optimally solve anything but the most straightforward problems. In other words, the human scheduler either has to almost continuously jump from one aspect of the problem to another or, as actually happens, ignore the repeated requirement to consider multiple aspects of the problem and make a non-optimum solution.

Job shop scheduling entails the construction of schedules meeting a number of constraints. Such constraints can be considered to come from separate sub-domains of the overall job-shop scheduling domain. For instance, the requirement for individual orders to be delivered on time can be considered

a particular specialised aspect of the scheduler's expertise. In some cases the nature of the order (or the customer) means that orders can be delivered later than the ideal delivery date. Similarly, some orders may be delivered early without the requirement to utilise valuable warehouse space until the customer is prepared to receive them. The scheduler must be expert at "constraint relaxations" of these kinds.

Another aspect of the scheduler's expertise concerns the selection of machines to undertake the processing or manufacture of any order. For instance, small screws are ideally made on a small screw-making machine. However, if it improves the feasibility of the schedule, the scheduler may use his knowledge about screw production to re-schedule the production of small screws on a machine ideally intended for the manufacturing of large size bolts and screws. Again this is an example of constraint relaxation. The constraint in question is that items should be made on machines optimised for the size of item in question. Constraint relaxation expertise is used to relax this "soft" constraint in the case of "small" items being made on "large" machines. Usually the reverse, where "large" items were to be made on "small" machines, would be physically impossible and constraint relaxation expertise could not relax this "hard" constraint.

A further sub-domain of the scheduler's expertise concerns the need to maintain the equipment of the job shop in good order and for this he must be able to schedule routine maintenance. Again there are constraints and, for instance, a machine must not be left unmaintained for too long as this will predispose to sudden mechanical failure. (This is exactly the sort of issue already considered in Engineering Preventative Maintenance.)

Another sub-domain of the scheduler's knowledge concerns lot splitting. To

what extent is it efficient to split an order so that it can be fitted in the gaps between other jobs. There is a tradeoff between the ability to produce goods and the extra effort required to set and re-set the tooling on the machine.

These are just four of the separate sub-domains which job-shop schedulers may consider. In any particular job-shop there will be a fixed set of sub-domains (perhaps five or six) which must be considered during the construction of a schedule. Most of the references to any sub-domain will concern constraint relaxation as provided in the examples above.

The inter-sub-domain references are usually so heavy that it is realistic to consider that co-operative problem solving is involved, calling upon various knowledge sources each relevant to a particular sub-domain. A blackboard architecture is potentially the most suitable way of representing such problem solving.

Nevertheless, when the job-shop is very much less than saturated, just one sub-domain search can satisfactorily complete the schedule without the requirement to refer to other sub-domains for constraint relaxation. It is only as saturation is approached (perhaps demonstrated by a bottleneck in the schedule) that constraints must be relaxed. After this point, a situation soon arises where the pressure of inter-sub-domain references exceeds that which will enable an optimal schedule to be approached.

Scheduling also illustrates the use of inadequate trial solutions iterated back into a new cycle of problem solving. This is similar to the situation found in Audit Work Planning. A scheduler can recognise the inadequacies or lack of optimality in a schedule which he has just produced and will use this

information as input to a new attempt to produce a better schedule.

Diagrammatic Representation of Problem Solving Components

In this diagram the various sub-domain components are more heavily linked and, as with Audit Work Planning, there is a clear iterative feedback loop from the conclusion stage back to the initial data acquisition.

Of course, the final schedule in all its complexity is the final "Grand Overarching Conclusion".

Multiple Simple (or Co-operative) Search with Heavy Reference Linking - Fixed Set of Sub-Domains

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Heavy Reference Linking

Iterative Feedback

Grand Overarching Conclusions
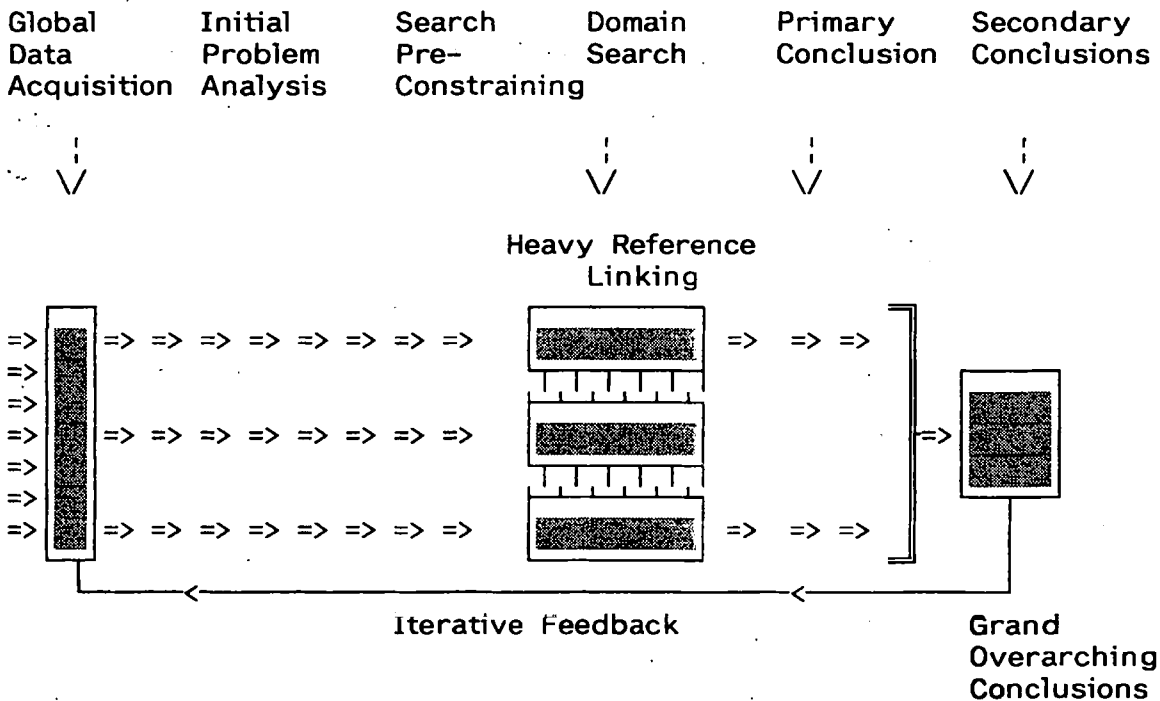
Commercial Implementation Details

This system has not yet been implemented due to lack of sponsor's budget.

However, other factors would have made implementation difficult and the system only partially effective. Key implementation difficulties would arise simply because of the complexity of the interactions between the various knowledge sub-domains.

Additionally, the system would have been only partially effective because downturn in business demand fundamentally altered the job shop constraints. At times of economic boom, job shops are working at nearly 100% of the theoretical capacity of the machines in the job shop. Staffing is not an issue since the workload can support the hiring of extra staff to man the machines. Therefore the task of the expert (and of the proposed expert system) was to schedule the machine usage to enable the maximum throughput for the job shop.

However, during recession, the job shop is not required to work at full capacity and a proportion of the staff are made redundant. In this case the problem becomes one of scheduling the staff so that they can work at nearly 100% capacity. Since equipment cannot be made redundant there are now excess machines and there is no real need to construct a complex schedule of machine usage equipment. In fact, the main requirement is to avoid setting-up effort on machines since this is a prime call on staff time.

The proposed system was first considered when the job shop was at full capacity. It did not undertake staff scheduling and thus was not really fitted to the problem when downturn in demand changed the constraints from machines to staffing.

Additional Comments

Two further aspects of the "Job Shop Scheduling" task are also important when considering the feasibility of applications to support this type of task. To clarify the first aspect it is considered with reference to analogous situations in holiday companies and financial trading. The second aspect is concerned with the consequences of relationships between a hierarchy of schedulers.

Job shop scheduling is a difficult task to support with an expert system but human experts make a reasonable job of it. They schedule "jobs", which are production runs of various manufactured items through a "job shop" consisting of a range of universal machines which can be set up in combinations to perform the various stages of manufacture for each job. In the main the experts produce a "satisficed" schedule, which although not optimised, gives adequate use of the theoretical throughput capacity of the job shop. (In an ideal schedule, jobs would be scheduled to use 100% of the job shop's capacity with no machines "waiting" for another machine to finish its stage in the manufacture of a particular job.)

One reason why a human scheduler may succeed where an expert system fails is that the job shop scheduling problem is not fully described in advance of the requirement to produce a schedule. In consequence the scheduler must make and manage relationships with the customers. To make explanation easier this situation will be described by using the analogy of a holiday tour company.

The holiday tour company pre-books blocks of hotel rooms and air flights and coach journeys between tour venues in advance of the season. (These

192

can be considered as equivalent to installed machine capacity.) Intending holidaymakers send in application forms for the tour itineries that they require and the chosen weeks of their holidays. (This is roughly equivalent to the "jobs".) If bookings are poor for the season then it is easy for the holiday company to produce a schedule which gives everybody the rooms and journeys that they require. But the holiday company makes a loss since there are many vacancies in the schedule.

If bookings approach the theoretical capacity then it is inevitable that a workable schedule cannot be created since individual preferences will inevitably mean overbooking of some rooms at certain periods and underbooking at others. (This is equivalent to customers asking for their jobs to be completed at due dates which happen to coincide to the extent that a "bottleneck", where a particular machine is overbooked, is created.)

To ensure that the holiday company can satisfy the maximum number of customer bookings and maximise its profits, the company will telephone individual customers to ask whether they will accept holidays at different dates or locations. Usually, although they have booked specific dates and locations, customers will have a degree of flexibility and a sufficient number of customers will be able to make adjustments to their dates and locations to enable the company to prepare a viable and profitable schedule.

The same negotiation occurs in job shop scheduling. When a bottleneck is apparent the human scheduler asks customers whether, for instance, they will accept early delivery (which is usually possible if they have sufficient stock space) or whether a small delay in delivering the order, or part of it, will cause a problem. By making a series of telephone calls the human scheduler can usually adjust jobs around the bottleneck so as to effectively

remove it and so produce a viable schedule.

This process is very difficult to reproduce with an expert system because, in essence, the system needs a human assistant to telephone and negotiate with customers during the production of the schedule. It is not possible to pre-specify the flexibility of each customer (ie the problem constraints) to enable a viable schedule to be produced without reference to the customers during the scheduling process.

If such a "scheduler's assistant" is necessary to assist any expert system scheduler then it is difficult to persuade the sponsor that the expert system is competent and cost-effective. Similarly, and a more insidious problem, is the fact that the human scheduler is also an expert on the management of customers. For instance, he will know (or contact his Accounts Department to check) that a particular customer is late with his payments. In such cases it is more likely that the customer will accept late deliveries or partial deliveries with a second delivery on an alternative date. In other cases, there will be customers who have been in dispute with the company over product quality or failure and who would be very annoyed with a suggestion that the delivery date of their current order should be changed. In such cases the scheduler will not attempt to call the customer to suggest a change in their order.

There will be a range of other circumstances where the human scheduler uses specific knowledge to manage the customers and it is difficult either to encode these or use them within an expert system. (An analogous problem occurs in the relationships between financial traders who have "he owes me a favour" relationships. One trader will often get a special low price from another trader because he provided a similar favour to the other

194

trader previously. One prototype financial trading expert system actually provided an "owes me one" flag (with two levels of value) for traders to record their debts and owings to other traders.)

The second aspect of feasibility of Job Shop Scheduling task arises because scheduling is actually conducted as a hierarchy of schedules. The master scheduler (together with other experts as described above) prepares a monthly rolling schedule (as a Gantt chart) for the production of the orders of various types of products that month. Orders are distributed throughout the month and are for products such as television sets in a range of screen sizes, cabinets and combinations of facilities such as stereo sound, teletext, remote control, vision-in-vision, UK and other European signal standards, etc.

The master schedule indicates activities such as "component insertion" on specific printed circuit boards (PCBs), final assembly into cabinets, etc for each of the orders. As described above the master schedule is actually prepared on special whiteboards, in the form of large Gantt charts, covering the walls of the scheduling room. Once this schedule has been produced it is transcribed onto paper and distributed to the various departments of the factory. Here it becomes the basis of orders to those departments. For instance the PCB department must cut large sheets of copper plated glass fibre sheets into the sizes for each board, etch the circuit tracks, drill the component holes and clean the boards. The cabinet department must prepare the resins and extrusion presses so as to produce cabinet components for assembly into cabinets to house the TV chassis.

The result is that the other departments also have to produce schedules to satisfy the orders from the master scheduler. In some cases the sub-

schedules produced by these departments become orders to further departments (or external supply contractors) and there is thus a third layer of scheduling to be undertaken. For example, certain printed circuit board sub-assemblies are "components" of the TV master printed circuit board. The production of such sub-assemblies demands a sub-schedule. However, the machines which insert the components into the printed circuit are reliant on "bandoliers" which are like machine-gun bandoliers but consist of repeated sequences of components mounted on twin paper tapes and in the right order for insertion. The bandoliers must be produced by a specialist contractor who needs to schedule this work as part of his overall load. (The production of components could be considered as a fourth sub-layer of the schedule but manufacturers hold sufficient buffer stocks for this to be irrelevant.)

The scheduling hierarchy works well in normal circumstances. The master schedule takes about a week to prepare and the sub-schedules each take about three working days. Provided the schedules are commenced a month in advance this causes no problems. However, if there is an unexpected manufacturing problem, the process is thrown into chaos.

An extreme example occurred when a new variant of a television set was being produced. A new tuner had been designed and this was to be incorporated in the television sets. However, when the first sets reached the end of the production line it was found that they failed to work properly. The tuner had been designed outside the UK and had been tested against simulated UK PAL broadcasts. However, the simulated broadcasts were incorrect and tuners did not work against actual signals in the UK.

The production had to be immediately halted and a new schedule produced

to enable orders of proven models to be completed. However, even by working on an emergency basis, the new schedule took about four days to complete. The sub-schedules for the PCBs and cabinets also took a couple of days with the result that the factory was at a standstill for a working week awaiting a schedule. The financial loss was enormous because the large number of machine operatives had nothing useful to do.

Had an efficient expert system scheduling system been in place rescheduling at the various levels could have been undertaken in a matter of hours. In fact, hierarchical scheduling would not have been necessary since a single system would be able to undertake all the scheduling levels. Nevertheless, the problems of managing scheduling relationships prevented a full implementation so that the factory still uses human schedulers with the risk of catastrophic stoppages remaining. Even though, for the reasons described above, any scheduling system would have required an expert human assistant, we suggest that the availability of such a system to reschedule in catastrophes would in itself justify the implementation investment.

3.14   TELECOM REPAIR SERVICE

Sponsoring Organisation

This investigation was sponsored by British Telecom and was conducted as part of the author's work with that company.

Introduction

The investigation of the Telecom Repair Service task formed the initial basis of this research. In the event it caused a major change of career for the author and in turn this career transition markedly affected the course of the research.   As a result the Telecom Repair Service investigation report differs to and is more detailed than the remainder of the reports in this chapter.

Description of the Telecom Repair Service Task Domain

The Repair Control Officer (RCO) is a skilled telecommunications technician employed in a British Telecom Repair Service Centre (RSC).   Such RSCs are the focal points for the reception, diagnosis and progress of repair of telephone faults in a local telephone network.   Within this working environment the RCO, together with his colleagues, may have to deal with the diagnosis and management of faults on any one of up to 100,000 telephone lines served by up to twenty telephone exchanges.

199

The RCO, who usually has a background of considerable "field" experience, works at a "Test Desk" which is similar in appearance to the operator's "switchboard" as used in the period of manual telephone exchanges. The Test Desk is equipped both with incoming "complaint" lines from customers and also with special circuits, routed via the served telephone exchanges, which can be used to make electrical tests on customers' lines.

The RCO receives telephoned fault reports from customers, makes diagnoses of the likely cause and location of faults and acts as a field staff controller to progress clearance of faults. Our interest lies in the RCO's diagnostic role including both electrical testing, diagnostic questioning of the customer and the use and interpretation of records.

It was known that many exchange line diagnoses were poorly made and that there was considerable performance variation between individual RCOs. In the United Kingdom decisions made by the subjects of our investigation commit the activities of some 20,000 faultsmen and decisions to unnecessarily or inappropriately commit specialist staff are, in overall terms, very costly.

Reports of faults or poor service on customer's telephone lines are usually received by a Fault Receptionist (who may sometimes be the normal telephone operator) and passed to an RCO for diagnosis. The fault reports are either written on a paper "docket" or entered on a fault database. The diagnostic part of the RCO's job starts when he has before him the docket or screen detailing the customer's complaint together with a Fault Record Card (or its database equivalent) with details of the complainant's line. The latter document records permanent information about the line, eg customer's address, equipment installed etc, and also a "history" of previous faults.

We are concerned with that part of the RCO's job where he must progress from the fault report to a decision on action required. The range of options is wide. For instance, at one extreme, it may be decided that no further action is required (eg where the customer report is suspect or the fault untraceable). At the other extreme, a decision may be made to alert a cable gang to undertake highway excavation or other major work.

The RCO commences his diagnosis on the basis of information on the report docket and Fault Record Card. This information, combined with the results of electrical tests performed on the Test Desk and replies to diagnostic questioning, is used in conjunction with the RCO's own knowledge of the local network and general experience of fault diagnosis.

Knowledge Acquisition in the Normal Working Environment

As will be explained in Chapter 4, most academic research into problem solving, as part of the understanding of human cognition and latterly with the aid of computer simulation, has taken place in the laboratory or other place remote from the working environment.

We could not adopt this approach for Telecom Repair Service. As detailed below, diagnostic protocols were collected and there would have been considerable difficulty and expense resulting from the simulation of the "live traffic" (ie real faults from genuine customers) on which the RCO depends to demonstrate his skills. More importantly this task requires the use of a complex Test Desk coupled to special testing apparatus forming an integral part of a real telephone exchange. There was thus no alternative but to conduct the line diagnosis knowledge acquisition in the normal Repair

Service Centre working environment.

Use of Real Diagnostic Examples and Concurrent Verbalisation

Although the requirement to undertake knowledge acquisition "on site" could not be avoided for Telecom Repair Service, the choice of the actual method of knowledge acquisition was not constrained in the same way. Thus, although it was essential to undertake knowledge acquisition at the Test Desk, we were free to use any particular method of acquiring the RCO's knowledge. Nevertheless, in the event, the actual method chosen was intimately linked to the working Test Desk environment.

For Telecom Repair Service we chose to collect actual diagnostic protocols supplemented by the responses to supplementary questions from the experimenter. Apart from Eurobond Trading which involved protocols "by review", the method of knowledge acquisition used for Telecom Repair Service is rather at variance to methods used in all of the other tasks investigated. Such other tasks involved knowledge acquisition generally following a "conversational" style with frequent use of hypothetical examples etc.

Further, in some of the other tasks such as Benchtop Electronic Diagnosis we have sometimes used concurrent verbalisation where we have asked the expert to add a commentary to the task as he performs it. This commentary is made solely for the experimenter's benefit and, of course, is not made by the unobserved expert.

However, in Telecom Repair Service, the verbalisation (eg speaking to the

202

remote customer) is often an essential and already existing part of the task itself. In Telecom Repair Service a major part of the knowledge was acquired by recording this verbalisation and observing and recording the physical actions of the RCO in performing actual diagnoses.

Thus the investigation of Telecom Repair Service differs from the investigation of the other tasks reported in this research (except for Eurobond Trading) in that it both makes far greater use of actual (rather than hypothetical) examples and it also employs a larger element of knowledge acquisition by "observation" (both visual and auditory) rather than by direct interaction between investigator and subject. The reasons for these differing emphases are considered below.

Given time and opportunity there is probably no better way of sampling an expert's knowledge than by observing the natural run of his work over an extended period with the opportunity to question him as and where necessary. Unfortunately this process does not make particularly effective use of the investigator's time for the following two reasons.

First, unless the working examples (eg diagnostic examples) are carefully chosen, the distribution of cases appearing before the problem solver will inevitably contain a large proportion of common or routine cases. Thus the rare and more interesting conditions will only appear between the mass of repeated familiar cases. As a consequence much of the experimenter's time will be wasted on the observation of duplicate cases.

Second, in case of diagnostic/engineering tasks in particular, experts who make physical tests on their domain spend a significant proportion of their time actually preparing for and making the tests. Such ancillary information

gathering work can actually take far more time than the diagnosis itself. Thus in Benchtop Electronic Diagnosis the expert diagnostician often has a good deal of disassembly or unsoldering of equipment. Such equipment is usually designed for efficient operation rather than ease of fault diagnosis or knowledge acquisition.

Telecom Repair Service Rapid Test Oriented Environment

The RCO is in a very much more favourable position in that his diagnosis is more-or-less concurrent with his "testing" (eg actual electrical testing or diagnostic questioning of the customer). The whole apparatus of the Test Desk is designed for rapid testing of faulty lines.

Similarly, customer response to questions is also rapid without the need for preparations. Further, records are either available at a glance or can be rapidly accessed on a screen. Thus a far greater proportion of the RCO's time is devoted to actual diagnosis than, for instance, the diagnostician undertaking Benchtop Electronic Diagnosis who is obliged to carry out a range of preparatory or subsidiary actions. Therefore the second problem discussed above does not really arise.

Rapid "Unverbalised" Diagnoses

Both the RCO and benchtop repair diagnostician see the same high proportion of run-of-the-mill cases. In the case of Benchtop Electronic Diagnosis this potentially time wasting problem was avoided by dealing largely with hypothetical examples coupled with questioning about isolated

diagnostic techniques.

In Telecom Repair Service, and as explained above, it was not easy to use hypothetical examples. However, the investigations were relatively efficient even though an approach biased toward protocol collection was taken. This efficiency resulted as a consequence of the considerable difference in time required for an "unverbalised diagnosis" (ie verbalisation over and above that necessary for the task) and a diagnosis where concurrent verbalisation is adopted with the investigator exploring the verbalised diagnoses with supplementary questions.

Ericsson and Simon (1980) show that slowing of performance is the principle effect of concurrent verbalisation. This finding has been largely confirmed by Russo, Johnson & Stephens (1989) although they notice some slight loss in mathematical task accuracy when subjects are asked to verbalise. Nevertheless, they also note that retrospective protocols contain evidence of substantial forgetting and fabrication which they assert makes them much less valuable than concurrent verbalisation. Slowing of performance was the only observable and very marked effect with RCOs. Typically an RCO might deal with a "routine" diagnosis in less than two minutes. However, an "interesting" case can take up to fifteen minutes (or sometimes even more) where the technician has to explain his diagnosis and answer the investigator's supplementary questions.

Knowledge Capture

Before proceeding to the precise way in which diagnostic technicians were assisted to verbalise and explain their reasoning (ie the knowledge

205

elicitation process) we must first consider the knowledge capture process by which the elicited knowledge was recorded. Although we now prefer to use video recording as the medium for knowledge capture where the expert is performing electronic tests, observing meters etc, this was not possible at the time that Telecom Repair Service was investigated and so audio recording was adopted instead.

It is often suggested that knowledge engineers should make notes to record knowledge acquisition sessions with experts. However, we have never found it possible to take useful notes and at the same time manage the knowledge elicitation process. The social situation is inevitably somewhat unnatural and it seems far less disturbing to set up recording apparatus than to attempt to note down the subject's thoughts.

More importantly note taking is at best a selective process and it is not until sometimes well into the knowledge analysis stage that the significance of a problem solver's comment is realised. In many cases a vital clue to understanding has occurred at the mumbled end of a sentence and would almost certainly have been overlooked when note taking.

Tape Recording Procedures

Verbal data was collected using a small self-contained cassette tape recorder placed on the Test Desk where the RCO performs his tests and diagnoses. Subjects were given a short but careful assurance that their identities would not be revealed in any subsequent report and this, perhaps coupled with the unobtrusive nature of the tape recorder, gave rise to no instances where we were aware of any subject reticence due to the tape recorder.

In fact, in Telecom Repair Service or any other investigations, we have never experienced any problems with using a properly introduced tape recorder.

Verbal Data/Protocol Collection

Our method of verbal data collection was based on three pilot sessions, each of about an hour's duration. These sessions were conducted with a close working colleague of the author, experienced in fault diagnosis, who acted as subject. It soon became obvious, even with this co-operative subject, that any attempt by the author to impose a formal structure (as would be expected in laboratory-based research on problem solving) was very difficult in the early knowledge elicitation sessions. Also, when attempted, it had a very marked and disruptive effect on the subject's normal method of working. Nevertheless, as described below, later sessions with a subject did have an increasing structure based on questions ("errors and omissions") arising from the analysis of earlier sessions.

It was thus decided to accept that the course of each of the data collection sessions with a new subject (which might be of several hours in duration) would have a large proportion of unstructured content and it would be for the author to disentangle the messy data after the experimental session had been concluded. The verbal data consisted of problem solving protocols (perhaps involving dialogues with customers) together with a record of the conversations between the author and the subject. Since most diagnoses take less than ten minutes to complete it was possible to collect a wide variety of diagnoses in an experimental session.

Choice of Subjects and Method of Working

For the purposes of the Telecom Repair Service investigation, verbal data was collected separately from three highly skilled RCOs. The experts were selected on the basis of a mix of supervisor and peer recommendation. These three subjects, of whom one (Jim Harvey) provided the cornerstone of our data, were observed in a series of sessions, each of which lasted up to six hours and were tape-recorded throughout.

In all cases the author sat beside the subject at the Test Desk and, in a conversational mode, sought to clarify and explore the running comments made by the subject.

Hypothetical Cases closely related to Current Real Cases

Apart from their commentary on ongoing diagnoses, subjects were encouraged to make additional observations on hypothetical cases nearly similar to those with which they were currently dealing. Usually subjects were only sidetracked in this way after a particular diagnosis was complete and only in those cases where the author clearly recognised that the subject had made a clear EITHER/OR decision on the basis of a test result or customer's answer to a diagnostic question.

In these instances the author would ask the subject what would have happened in the alternative condition (ie alternative test result or answer). By this means it was often possible to elicit a hypothetical but confident (and therefore taken as valid) protocol for the "alternative condition" fault. This technique is particularly useful for gathering data on low-frequency

208

occurrences which, although familiar to the subject, might never be recorded within the limited data collection periods.

Agenda of "Errors and Omissions" Explored at Next Session

Somewhat similarly to the above, analysis of the protocols between experimental sessions highlighted a limited number of discrepancies, doubts and unexplored alternatives which were explored and recorded with the subjects during the next data collection session. In second and subsequent sessions it was normal to open the session by working through a prepared agenda of matters for resolution from the previous session. As a result later sessions have an increasing content of structured knowledge elicitation and a declining content of knowledge elicited as the result of verbalisation of current cases.

Law of Diminishing Returns

Knowledge elicitation followed a law of diminishing returns so that, from experimental session to experimental session, there was a declining probability that new fault circumstances and diagnoses would be encountered. For this reason the transition to structured knowledge elicitation must be eventually made so that, for instance, low-probability explorations of hypothetical conditions can be made.

Knowledge Organisation for Telecom Repair Service

Telecom Repair Service was unusual in that it resulted in far more verbal protocols than tasks subsequently investigated. Thus development of a method for the treatment of this potentially voluminous material was particularly important. Although later task investigations resulted in less direct protocol material, the experience gained in Telecom Repair Service has proved very useful in the handling of these subsequent investigations.

In the investigation of Telecom Repair Service the author decided that the attempted simulation of the RCO's expertise would run parallel to the knowledge acquisition sessions. However, we had no expert systems shell implementation skills at the early stage of this research. Further, commercial expert system shells, which might have served as prototypes for an implementation, were then not easily available. We did not wish to delay the Telecom Repair Service knowledge acquisition until a shell became available so, still wishing to give some test to the acquired material, we decided that our early "implementation" efforts would be on a "paper based" basis. Appendix A (Telecom Repair Service: Analysis of Acquired Knowledge) provides a sizeable extract of the various paper-based documents progressing from verbal protocol transcripts through to initial rule-sets.

As with the experience gained with handling extensive verbal protocols, the experience of knowledge analysis and structuring towards a "paper based implementation", although not repeated in later tasks, has been useful in our later work.

Verbal Protocol Transcripts

Transcripts were produced which contained all the recognisable information in the protocol but did not record the more or less redundant linguistic content found in normal conversation and explanation. Line numbering was used to aid later reference to the text. The further distillation of the taped protocols resulted in the exclusion of much that was irrelevant or was concerned with "management" aspects of the RCO's job. Appendix A includes extract examples of the original, distilled and analysed knowledge.

Transcript Annotation

Protocol were not only structured by the use of numbers, but fault diagnostic sequences were alpha-numerically labelled and their content distinguished from the rest of the text by tabulation. This results in a protocol which automatically divides itself into two classes of information. First there are fault-specific comments and actions appropriate to each labelled fault and, second, there is the remainder of the subject's commentary consisting largely of heuristic comments.

RCOs often deal with the diagnosis of faults on an interleaved basis so that several diagnoses may be simultaneously in progress. This method of working is occasioned by the fact that a particular diagnosis may be held up for want of consulting a record away from the Test Desk or by, say, the allegedly faulty line being temporarily in use or not testable due to the use of certain test circuits by another RCO. The interleaving of diagnoses results in a protocol where the complete record of a particular diagnosis may be scattered over two or three pages of the transcript.

Problem Solving Sequences

As a means of assembling disseminated diagnoses, problem-solving sequences were thus extracted so that each sequence showed the stages in each diagnosis uncluttered with extraneous comments and segments of other diagnoses. Examples of such sequences are also provided at Appendix A.

Where a fault sequence replicated that of an earlier fault, a new sequence was not prepared but the original sequence was annotated with the identification of the duplicate fault. The more common fault sequences thus became recognisable by their considerable number of annotations.

Heuristic Fragments and Strategic Knowledge

The problem solving sequences were used as the basis of rule compilation as discussed below. The remainder of the protocol, of which much we term as "heuristic fragments" (ie isolated elements of problem-solving knowledge), could not be directly assembled into any meaningful sequences. This material was thus used directly, by scrutiny, for the extraction of rules. However, even when the complete fault sequences and recognisable rules had been extracted, the protocols still contained a considerable amount of problem solving knowledge. In this residual material we were, for instance, able to isolate knowledge relating to the higher-level strategies used by the RCO. Much of this knowledge was used in the design of the inference mechanism of the Diagnostic Inference System (DIS). The acronym "DIS" relates to the diagnosticians colloquial term "dis" meaning a circuit disconnection.

212

The Diagnostic Inference System (DIS) – Choice of Knowledge Representation

At an early stage in the Telecom Repair Service investigation we concluded that a variant of production rules would be the most appropriate knowledge representation principally because it seemed best able to directly allow us capture the modular "chunks" of knowledge which we were eliciting from our subjects. Thus, in the main, our preference for this form of representation was made on the grounds of ease of knowledge transfer from human to system rather than on any argument in favour of the theoretical appropriateness of one representation over another. Anderson's (1976) arguments on the non-viability of a unique cognitive model and the equivalence of alternative representation-process repairs seemed to support our decision to opt for convenience and availability when choosing a knowledge representation.

Problem Solving Architecture

Telecom Repair Service differs from the tasks so far considered in that, although the task expertise can be divided into sub-domains (Report Interpretation, Diagnostic Testing etc), the final diagnosis can often be established independently from more than one sub-domain.

To clarify this point it is useful to refer back to Pensions Management considered earlier. Pensions Management has a division of the domain knowledge into sub-domains (Credits, Widows, Children etc) but these are each genuinely and uniquely able to answer enquiry problems in the sub-domain to which they are appropriate.

Thus although, as just discussed, the "Children" sub-domain must be referenced to provide an answer for the "Widows" sub-domain, it has no alternative competence to answer the "Widows" enquiry itself. Similarly, the Pensions Officer has no alternative but to use his "Credits" knowledge in order to determine if a pension scheme member is entitled to benefit from service in a previous company. His knowledge about the rights of "Widows" has no relevance to this issue.

In contrast, in the case of Telecom Repair Service both the "Report" and "Test" sub-domains can often provide equally competent diagnoses (ie problem solutions) to the same problem. Therefore, in many cases, it is possible to determine that the location of a fault is in a customer's premises (rather than in the street) either by detailed questioning of the customer ("Does it occur on all extensions?") or by testing the impedance of the line (to determine the number of extensions present). On the other hand, the "History" and "Records" sub-domains can rarely provide alternative independent complete diagnoses even though their domain coverage is equivalent to that of the "Report" and "Test" sub-domains.

As a result of the situation just described the inter-sub-domain references involve both support from sub-domain to sub-domain (as in Small Business Guidance, Job Shop Scheduling, Eurobond Trading and Pensions Management) but also, and more importantly, the strengthening (or "confirmation") of hypotheses. In Telecom Repair Service there can be no absolute confirmation of a diagnostic hypothesis until a faultsman has been dispatched to actually physically confirm and clear a defective cable, connection etc. Since such final confirmation is manpower expensive, a hypothesis generated by, say, diagnostic testing is strengthened by further customer questioning or analysis of the Fault History to determine whether

214

the hypothesis is plausible.

However, although other sub-domains provide hypothesis strengthening, the diagnosis is usually considered the product of a particular sub-domain. Usually this is the "Test" sub-domain since this represents potentially the most powerful of the parallel routes to a diagnosis.

If the diagnosis is proven to be wrong as a result of the faultsman's visit, and the fault still persists, this information is fed back on an iterative basis as additional information to a new diagnostic cycle. This is equivalent to the same aspect of problem solving seen in Audit Work Planning and Job Shop Scheduling.

Diagrammatic Representation of Problem Solving Components

The diagram overleaf shows that the domains are overlapping in their scope and that there is a single secondary conclusion (the location of the fault) which, if doubted by the diagnostician, is fed back as part of the initial data in a re-consideration of the diagnosis.

Co-operative Search with Reference Hypothesis Strengthening

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining	Domain Search	Primary Conclusion	Secondary Conclusions

Overlapping Domains with
Heavy Hypothesis Strengthening
Reference Linking

Iterative Feedback
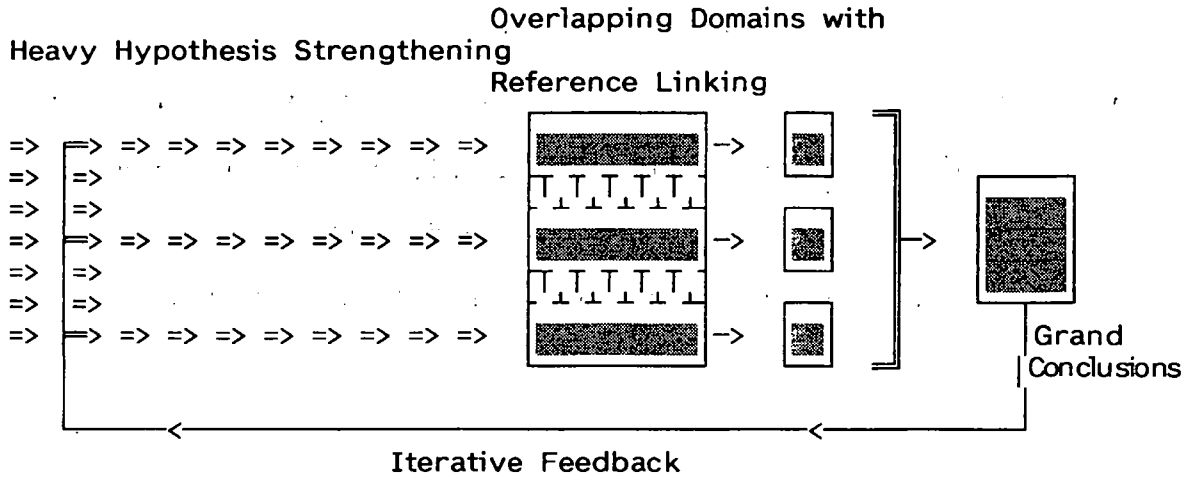
Commercial Implementation Details

The Telecom Repair Service system was potentially a "winner" but could not be progressed beyond a pilot for two reasons. First, the customer records and electronic line diagnostic systems (on which both the human diagnostician and the expert system were dependent on data) were implemented on a Honeywell minicomputer which could not support a full version of the Lisp-based DIS expert system. Second, concurrent improvements in the records systems (previously paper based) and improvements in repair service staffing gave such large potential for increases in efficiency and service performance that management were quite unwilling to consider the additional benefits that expert systems might

216

bring. With strong management support it is possible that the hardware problems could have been overcome. However, with key managers already very satisfied with their existing success, they saw no reason to invest for even further improvements by operationally implementing this system.

3.15   EUROBOND WARRANT TRADING

Sponsoring Organisation

This investigation was made in co-operation with a major international bank.

Introduction

Financial trading applications in the City of London are a favourite target for suggested AI/expert systems support.   Yet, although there are some reputedly secret successful systems, no widespread use of the technology has occurred.

For reasons of business caution, the author had avoided financial trading applications in favour of other financial sector tasks, such as Life Underwriting and Pensions Management, where there seemed a better chance of actually completing a useful expert system.

Nevertheless, an opportunity arose to speculatively investigate the problem solving of a leading Eurobond Warrant trader.   The method of investigation which was devised proved very effective and an understanding of the structure of the trader's problem solving was attained.   However, perhaps as expected, no feasible route to effect an implementation was originated.

Author's Role

Apart from some interpretation of specialist terms during knowledge

acquisition, where a colleague of the expert gave assistance, the design of the investigation, its execution and subsequent analysis were solely conducted by the author.

Outline Task Description

Financial traders, whether they are trading in financial futures, gilts or options have a key objective to set a price on the instruments in which they trade. In the case of the Eurobond warrant trader the responsibility was to set a price on some one hundred and fifty warrants which are separately tradeable from the Eurobonds to which they are normally attached.

The market can be quite volatile because the price of the warrants, since they imply a right to purchase the more expensive Eurobond, demonstrate disproportionate effects as a result of changes in the underlying market. Further, at the time this investigation was undertaken, the Eurobond Warrant market was quite new and the relative infrequency of trading any particular instrument gave very few benchmarks (ie actual trades) upon which the trader could base his prices.

Form of Investigation

The informal project solely involved an investigation of the trader's problem solving and subsequent analysis. One short introductory knowledge acquisition session followed by two one-day in-depth sessions were conducted.

The trading environment does not easily lend itself to knowledge acquisition. Traders are constantly fully occupied either on the telephone, using one of the many screen-based information sources or talking to colleagues. The only time that their activity seems to be suspended is when the market is moving to their disadvantage. At these times traders are often in the process of losing considerable money and, at this stage, they become very quiet and are not interested in cooperating in knowledge acquisition activities.

Findings on Nature of Problem Solving

In some of the other tasks investigated there has been only a limited opportunity to undertake knowledge acquisition yet, even though an expert system has not ensued, we have been confident that we have acquired a good understanding of the way in which the problem solving is undertaken.

In this task, Eurobond trading, we have had a worthwhile opportunity to study the expert's problem solving but cannot claim that we fully understand how traders trade. Nevertheless, this description and the suggested problem solving architecture discussed below cover the structure of the problem solving process and discuss those detailed areas where further investigation is required.

Problem Structuring Props (Tools and Aids)

The principal problem solving tools are simple spreadsheets and decision support systems. The trader needs to explore a number of alternative

hypotheses which might result from a proposed trading opportunity. There are now many rather sophisticated modelling systems available to the trader and, in the trading of Eurobond warrants, these were used extensively. In the task investigated these were primarily PC-based tools with no interconnection to data feeds etc.

The other "workspace" supporting the trader is made up of the many pages available from Reuters etc data feeds. The use of this "workspace prop" is considered further at the commentary at the end of this section.

Generic Data

A mass of data confronts the trader and most of this is not specific to the current trading opportunity or problem. Some data is electronic such as that arising from the thousands of pages on Reuters and Telerate (including the text pages such as News etc) and the remainder is paper based from newspapers and journals. There is also a mass of historical data on previous Eurobond warrant trades undertaken by the bank.

Specific Data

Specific data on Eurobond prices is provided by particular Reuters pages. Nevertheless even this may sometimes be regarded as generic data since, when the trader is acting proactively, this material will be scanned until the trader perceives an opportunity for trading. At such stage, as a particular potential trade becomes the focus of the trader's attention, the data can be regarded as specific data to the current problem. Changes in the data (ie

222

changes in prices during the period that a trade is anticipated) are purely problem-specific data.

Knowledge Sources

Multiple knowledge sources can be considered.   The following is an incomplete set:

"Feeds"

Knowledge involved in the interpretation of screen based data can actually be considered as involving two separate knowledge sources:

a)   Interpretation of screen-based numerical data, particularly prices.

b)   Interpretation of text-based background news data.

"Corporate Reports"

Knowledge used for the interpretation of the specific information, including Reports and Accounts, of the company issuing the Eurobond to which the traded warrant is attached.   This knowledge source is similar, although at a higher level, to that employed by the experts for Small Business Guidance in that the trader must attempt to ascertain, very quickly, the "health" of the target company from published information so as to judge the prospects for change in the bond price.

"Counterparty"

Traders establish market prices, especially for infrequently traded instruments, on the basis of telephone requests for bids and quotes from market counterparties (ie other traders trading the same instruments).

Knowledge enabling the conducting of such conversations and the interpretation of subtle clues derived from them can be firmly considered as derived from a particular knowledge source. There are similarities to the "Questioning" knowledge source in Telecom Repair Service which- supports diagnosis on the basis of customer questioning.

Exposure

Exposure management of the trader's position is an ongoing background task although but influences the risk that the trader will take on any potential trade. This knowledge source can be considered to be used before a trade is made to determine the overall exposure effect on the trader's position.

"Portfolio Screening"

An important knowledge source is involved in the routine serial searching of the trader's portfolio. The trader was regularly observed scanning through his portfolio of warrants to determine

224

which might be the next best to trade. Sometimes use was made of a spreadsheet or simple decision support tool in an attempt to predict the effect of a changing market on a particular interest. In other cases it was just a routine sequential scan which the expert explained was just used for "memory jogging".

If the reminder about the holding of a particular warrant coincided with some new knowledge about the world (eg the coincidence of a motor manufacturer bond held at the time of impending industrial unrest) the trader might decide to change the price of the warrant. For instance, in the motor manufacturer example just quoted, the industrial news might induce the trader to lower the price of the warrant in an attempt to sell before a major dispute.

The knowledge involved in portfolio screening calls heavily on other knowledge sources. However, interestingly, portfolio screening is often the start of proactive problem solving.

"Extrapolation"

This term is loosely used for the knowledge involved in either deriving a price based on a known trade some days earlier or deriving a price based on the current known price of analogous bonds. Where instruments are infrequently traded traders have developed skills to enable them to quote a price based on such previous or allied prices. In some ways such knowledge offers a sort of in-built decision support tool except that it is based more on analogy than mathematics.

Interpretation of Salesmen's Feedback

In the trading environment traders are supported by salesmen who deal primarily with non-counterparty customers. Such salesmen provide a regular feedback on market rumours and trends. Traders have developed considerable skills in interpreting the casual fragments of information that originate from the sales force. Again, like information from counterparties and general traders' gossip, the feedback from salesmen is verbal data and would require to be (expensively) codified if to be used by an expert system.

Open/Closed Problem

Eurobond warrant trading offer a prime example of an open problem. There is no closed set of solutions and, for instance, the trader has an infinite range of prices from which he can determine a particular quoted price as the result of his problem solving.

Proactive/Reactive Problem

Traders are reactive to particular needs to buy or sell or to provide a quotation to a market counterparty. Their activity is proactive as a result of their overall "Make Money" objective.

Problem Solving Architecture

Eurobond Warrant Trading is in many ways similar to Job Shop Scheduling just discussed. However there is one important difference. Whereas in the case of job-shop scheduling the considered sub-domains were constant, in financial trading, as exemplified by Eurobond Warrant Trading, the set of considered sub-domains is not constant.

There is a core set of sub-domains plus certain frequently or infrequently explored sub-domains dependent on the course of the problem solving. The core set of sub-domains includes such things as interpretation of information feeds (Reuters, Telerate etc), interpretation of corporate news and figures on the company issuing the warrant, interpretation of information from counterparties (eg their offered prices) and maintenance of the trader's overall exposure. In addition the trader may separately explore sub-domains interpreting world news, industrial news (eg strikes and pay settlements), foreign exchange figures etc. The decision to explore such other sub-domains depends on an initial assessment of the problem. For instance, if a bond warrant was attached to a car manufacturer threatened with an industrial dispute, then the industrial news section of the Financial Times would be very actively considered.

The fact that the trader must not only decide when to invoke a sub-domain for search or reference but must also cope with heavy referencing between the "fixed" sub-domains makes the trader's task combinatorially very difficult. Correspondingly "correct" (ie leading to overall profit) problem solutions are only attained for just over 50% of problems encountered.

Diagrammatic Representation of Problem Solving Components

The diagram on the next page shows Search Pre-Constraining, driven by the nature of the domain problem, to limit the search to just two (heavily shaded) sub-domains out of a variable number of domains that might be considered for any problem. The sub-domains are heavily reference interlinked and the primary conclusions are summated to provide a Grand Overarching Conclusion on what should be the price of a particular bond warrant. As in previous tasks the trader may be dissatisfied with the outcome of his own thinking. In such cases there is feedback to change the nature of the way the problem is constrained. Usually this involved the trader deciding to widen the scope of the considered sub-domains to involve, for instance, conversations with counterparties and colleagues. Following this the trader iterates through the problem again to come to a revised conclusion.

Multiple Simple (or Co-operative) Search with Heavy Reference Linking – Variable Set of Sub-Domains

Global Data Acquisition	Initial Problem Analysis	Search Pre-Constraining (Domain Driven)	Domain Search	Primary Conclusion	Secondary Conclusions

Variable Overlapping Domains with Heavy Reference Linking

Grand Overarching Conclusions

Global Feedback

Commercial Implementation Details

Although the bank was very hesitant about permitting this investigation, the expert was very keen and it was possible to undertake useful knowledge acquisition.

Three major problems prevented implementation. First, funding could not be obtained. Second, and more important, it was revealed that average experts are right in bond warrant pricing on only about 52% of occasions. Key experts may be right on about 55% of occasions. Therefore bond

229

warrant traders have much the same status as successful horse race punters. In most expert systems work, and in the consideration of any expert's problem solving, it is usually taken as a prerequisite that the expert solves problems "correctly" in the vast majority of cases. This seems almost to be an implicit condition within the definition of "expert".

As a result a worthwhile expert system could not be contemplated where the underlying knowledge was so "shaky". This problem characterises many dealing/trading applications in the banking sector and many reputations, jobs and much funding has been lost attempting to implement expert systems on such an unsound basis. In fact, efforts in the trading sector have partially spoiled the reputation of expert systems in banking. This is to be regretted since many applications in retail banking (ie routine loan application and similar problem solving) are very easy to implement and have great utility.

A third problem, coupled with the unreliable nature of the traders' knowledge is the fact that traders change their knowledge (or as they express it, their "view") on an ongoing basis. Thus, a set of circumstances tomorrow might be interpreted differently from the same set of circumstances today. Where the traders' knowledge is changing, or as they would euphemistically say "evolving", it is difficult to envisage an expert system that does not have the capacity to update its own knowledge base in line with the expert's current view of the world.

The fourth problem in the potential Eurobond application was that traders acquire considerable verbal data on traded warrants from other traders. Like all financial trading there is much inter-trader gossip and discussion of rumours. All this verbal material has an effect on the traders problem

solving and price setting. However, even if the trading knowledge had been stable, the requirement to input a mass of verbally acquired data would have effectively killed the chances of the expert system being useable in practice. Similarly, part of the input to decision making is an interpretation of non-electronically encoded data (where Financial Times reports are usually quoted as the primary example). The selective encoding of such input, which can be voluminous and ongoing, is a task that cannot be undertaken by the trader and presents a further considerable stumbling block to the successful introduction of expert trading systems.

A fifth issue is that, as discussed above, the trader's interpretation of various screens etc seems to rely on a form of "perceptual knowledge" which is not easily articulated or captured. This presents yet another potential difficulty on the route to a viable expert system.

Although it is technically feasible a final sixth difficulty is that there would be considerable interfacing overheads in enabling any expert system to select and draw upon the many data feeds and individual data sources used by the trader.

In summary, the last point emphasises the need for data required for problem solving to be in a convenient form and, ideally, already electronically encoded. If there is a barrier, such as the requirement to encode a mass of verbally acquired data, then the system stands little chance of successful user acceptance even if the knowledge is highly stable and reliable. The fact that Eurobond Trading knowledge is neither stable nor reliable merely compounds an already very poor chance of successful implementation or use.

Additional Comments

These comments follow particularly from the brief introduction made above to the use of the multiple pages from external data sources (eg Reuters and Telerate) which constitute a problem solving prop around which much of the problem solving is undertaken.

The trader would constantly select and scrutinise a great many pages in the course of his trading. Although he would report when a particular figure or report lead him to initiate a trade or seek some further information it was not clear why this particular figure should be uniquely significant.

In the knowledge acquisition sessions (which were conducted "by review" as discussed in Chapter 4) the trader could amply justify his actions. Nevertheless such justifications seemed always to be post hoc explanations rather than the reasons why the data item was selected. This key issue is returned to in Chapter 6 in a discussion of the role of perceptual knowledge.

Similarly, and as considered above in the discussion of the architecture of the problem solving, there are multiple linkages between multiple sub-domains. Although such overall structure is very clear the detailed reasons why particular sub-domains are invoked during problem solving are many and various. In fact, although the expert will always give a post hoc justification (similarly to above), there seems little hope of capturing reliable generic knowledge covering the referencing between sub-domains.

Thus, in overall conclusion, there is no insurmountable difficulty in acquiring and understanding the overall structure of the Eurobond warrant

232

trader's problem solving. Even though it is a fairly complex architectural structure, as described above, we are confident that it represents the way in which problem solving is undertaken.

However, and this seems to be the current barrier to representation and implementation, it is difficult if not impossible first to acquire the detailed knowledge of how traders perceive individual relevant fragments of data from a mass of inputs and, second, how they select from a variable set of sub-domain knowledge to pursue the problem to solution.

The handling of knowledge such as that used by traders presents special problems for the expert systems developer. Coupled with the other difficulties already outlined for this task, it is not difficult to see why so many brave attempts to provide expert trading systems have failed.

3.16   SUMMARY

This chapter has reported on a series of occupational problem solving tasks investigated to greater or lesser degrees as potential candidates for expert systems implementations.   The next three chapters utilise the findings reported in this chapter to discuss the techniques needed to investigate occupational problem solving , the problem solving components identified in the tasks investigated and the factors which can affect expert systems feasibility, success or failure.

3.16   SUMMARY

This chapter has reported on a series of occupational problem solving tasks investigated to greater or lesser degrees as potential candidates for expert systems implementations.   The next three chapters utilise the findings reported in this chapter to discuss the techniques needed to investigate occupational problem solving , the problem solving components identified in the tasks investigated and the factors which can affect expert systems feasibility, success or failure.

CHAPTER 4 - KNOWLEDGE ACQUISITION TECHNIQUES AND EXPERIENCE

4.1   INTRODUCTION

It has been an important objective of this research to understand the problem solving of occupational experts.   The author has also had a concurrent commercial objective to acquire commercial expertise for the purposes of expert systems implementations.   Such objectives have demanded the development and refinement of investigational techniques.   There have been many studies of problems solvers yet what has become known as "knowledge acquisition" for many workers still remains a critical bottleneck (Feigenbaum & McCorduck, 1984) in the construction of expert systems.   The author has experienced and, in some cases, overcome considerable problems in acquiring knowledge.   In consequence, this chapter reports on the more useful of those techniques which have assisted the investigations reported in this thesis and, in many cases, have actually been key enabling factors in implementing commercial expert systems.   Towards the end of the chapter is a special section on "Problem Solving Props" that have been encountered during knowledge acquisition.   The study of the various physical devices used by the expert to help him cope with his work provides valuable clues to his problem solving and, in many cases, offers an easy route to providing similar functionality in the target expert system.

4.2   BACKGROUND

AI research on problem solving has a long history and there are numerous examples (eg Newell & Simon 1972) where human "experts" have been

observed in their attempted solution to a problem. In many cases these observation are accompanied by the subject's commentary on his progress to a solution or are interrupted or later followed by an experimenter's interrogation of the subject in an attempt to determine the internal mechanics of his problem solving. In the applied expert systems field such activities are often labelled "knowledge acquisition" and, for instance, Welbank (1983) offers a review of earlier work and Neale (1988) provides a good overview of work in the 1980s together with an extensive bibliography. (Correspondingly, the use of software tools to support knowledge acquisition is discussed in, for instance, Gaines & Linster (1990) and Laublet et al (1992). Use of such tools is in principle very much supported by the author but, in practice, he has not to any extent yet had an opportunity to productively use them.)

In the fields of occupational psychology, work study and systems analysis there have been even more numerous studies, at all levels of competence, of workers tackling problems encountered in their jobs. The motive for such studies is usually to increase the efficiency of the problem solving, either, directly, by the education or re-direction of the problem solver, or by providing better supporting equipment or documentation or, indirectly, by increasing the problem solver's job satisfaction.

However, even with this long history of research and commercial investigation, knowledge acquisition still remains a difficult area for the expert systems implementer. The KADS initiative was set up with a major objective to overcome such difficulties. However, KADS was not available for the majority of the period that this research was conducted. Even so, and as a supporter of the KADS initiative, the author is not yet convinced that it has overcome all of the problems of knowledge acquisition in the

commercial context. (This issue is further discussed below and is more fully summarised in Chapter 7) Nevertheless, as part of precursor work to KADS, the work of Breuker & Wielinga (1984) provides yet another good review of knowledge acquisition techniques.

The author very much recognises the difficulties associated with knowledge acquisition and suggests the following as important reasons for them.

First, it is strange that so many psychology-based studies supporting expert system research are still centred on tasks which are no more than the exercise of commonsense on an artificial problem. Nobody, except the academic, earns a living as an "expert" on the Eight-Problem or the Tower of Brahma. This situation has resulted in knowledge acquisition for real-life problem solving situations being left in many cases to amateurs from other disciplines. The psychologist has too often applied his rigour to the amateur cryptarithmetician and left the oil-well driller to the computer scientist.

Controlled studies in the Psychology Laboratory have an important contribution to problem solving research. However, the conditions necessary for control are sometimes incompatible with the investigation of a real-life task and, under such circumstances, the author would choose to sacrifice rigour for the opportunity to discover how such a task was accomplished.

Second, there are severe consequences to "non-clinical" nature of the sites where many real problem solvers undertake their work. Knowledge acquisition can be relatively straightforward when it is possible to isolate the problem solver to the confines of the Psychological Laboratory or even to a quiet office. In some of the tasks reported in Chapter 3, such as

Multinational Tax Planning and Small Business Guidance, it was possible to divorce the problem solver from his day-to-day environment and conduct knowledge acquisition in a quiet room. However, choice of the knowledge acquisition site was not feasible for most of the research on engineering problem solvers (such as Telecom Repair Service or Engineering Preventative Maintenance) whose work required them to use specialised tools and equipment to complete the tests and operations that were an essential part of their problem solving.

Even in those cases where it can be accomplished, taking the expert to the empty laboratory endangers the credibility of the acquired expertise. For instance, in Benchtop Electronic Diagnosis, much of the expert's knowledge was demonstrated (using equipment) rather than verbalised. In this and similar cases (including Job Shop Scheduling where the working environment consisted of a room whose walls consisted of Gantt chart whiteboards, the expert required "domain prompting" (ie by looking at equipment etc) to be able to recall the detail of his work.

So, in summary, there are special problems concerned with the investigation of problem solvers undertaking real jobs in the working environment. It is surprising that, amongst the papers on knowledge acquisition, techniques for acquiring occupational knowledge "on-site" are not well reported.

4.3   SPECIFIC TECHNIQUES

Examples of the main techniques used by the author in the course of this research, together with references to relevant previous work, are introduced under the following "Specific Techniques" topic headings.   Comments on

other issues relevant to some of the techniques, such as the commercial pressures on knowledge acquisition and "perceptual knowledge", are interspersed at relevant points.

To a large extent the following techniques represent a personal "toolkit", along the lines of that proposed by Gammack and Young (1985), adapted and developed by the author on a needs driven basis over the course of this research. The author firmly agrees with Kidd (1987) that there is no single "magical technique" to overcome the knowledge acquisition bottleneck and that a range of techniques, related to the challenge of the tasks investigated and the personal aptitudes and skills of the researcher, must be adopted on an evolving basis.

Specific Technique: Briefing

Wherever possible the author always briefed himself on the task before commencing knowledge acquisition. This is partially to maintain credibility with the expert and partially to understand the "buzzwords" (as with Breuker & Wielinga, 1984), and background to the task and the domain. Briefing was sometimes quite extensive so, for instance, the author spent part-time of several weeks studying international tax legislation before meeting the senior expert for Multinational Tax Planning.

Ideally, as with Multinational Tax Planning, books or manuals could always be used for briefing. Nevertheless many applied expert systems projects are mounted precisely because good domain documentation does not exist and the primary or most current knowledge is retained by the expert. For instance, in the case of Eurobond Warrant Trading there was no accessible

information on the trading of Eurobond warrants. It was precisely the newness and novelty of trading the instrument that made the expert's knowledge so unique and valuable.

Where no documentary evidence existed, the author usually used a "briefing expert" who was a co-equal or colleague of the actual expert. This was done with Small Business Guidance and it was possible to "waste credibility" with the colleague in order to make a non-ignorant approach to the real expert. In other cases, such as Audit Work Planning, the managerial project sponsor was very successfully employed as a briefing expert. Such persons are aware of the circumstances and language of the domain even though they may lack detailed or, more commonly, up-to-date expertise. Often such higher level views are important to understand the overriding purpose of the task which the expert undertakes.

A good further example was provided by Life Underwriting where prior referral was made to a senior underwriter not involved with the day-to-day business of underwriting lives. His explanation was extremely valuable as preparation for knowledge acquisition with expert "front line" underwriters but, more importantly, provided information on the business priorities of the life company which were not appreciated at the junior level. As reported in Chapter 3, most life underwriters think that their main goal is to be careful in assessing risks; their managers know that the real goal is to underwrite the maximum amount of business. Under these circumstances the few bad risks will be submerged by the profits from the increased business. As a result of the management briefing both knowledge acquisition and eventual systems implementation were successfully skewed toward "turnover" than accuracy in risk assessment.

<u>Specific Technique: Traditional Verbal Protocol Collection</u>

Telecom Repair Service was the first task investigated in this research and it was decided that the subject's problem solving behaviour would be recorded by verbal data collection.

Since the seminal work of Newell and Simon (1972) on cryptarithmetic, chess etc, there has been a long tradition of verbal protocol collection and analysis in the investigation of human problem solving. Where appropriate, and sometimes as an alternative, verbal data has been supplemented by the record of overt behaviour where the problem solving investigated has involved some manipulation of domain objects by the subject. Blocks world and seriating task problems are examples where a record of the subject's actions is invaluable.

In the investigation of Telecom Repair Service, the fault diagnostician shows overt behaviour in diagnostic questioning to the customer, in the various switch selections and testing routines performed on the Test Desk and in the accessing of specific customer and equipment records. However the only tangible evidence of the course and results of his diagnosis is the notes and diagnostic code written on the docket. Although there is thus considerable overt activity by the diagnostician it is still difficult to "induce" the methods by which he solves problems. For this reason, and because it was initially impossible to investigate this task "off-line" away from the equipment, it was decided to use the method of concurrent verbalisation in an attempt to establish the course of his problem solving. Newell and Simon (1972) was used by the author to provide a methodological framework for the method of verbal data collection which we used for Telecom Repair Service. The methods used are reported more fully in the

Chapter 3 account of the investigation of this task.

From Concurrent Verbalisation to Interviews

Protocols from concurrent behaviour were not appropriate for most of the tasks and, including later stages of knowledge acquisition from Telecom Repair Service, we soon moved to an interview style of knowledge acquisition. As our experience progressed we used and developed the techniques described in the following pages to increase the overall efficiency and management of the knowledge acquisition process. Nevertheless, it is first useful to consider the special constraints necessary when investigating problem solving in the commercial context.

The Commercial Constraints and Pressures on Knowledge Acquisition

Our work on commercial expert systems has caused us to very carefully consider the efficiencies of the expert system implementation process and, very particularly, the knowledge acquisition process. In addition to all the other criteria for commercial success, expert systems must be built within the cost and expert availability/tolerance constraints of the sponsor. Expert systems sponsors commit both their money and their experts' time to an expert system project. The project will not be viable if either resource is squandered or abused.

In fact, several of the tasks reported in Chapter 3 have not resulted in viable expert systems. Of course, in some cases, there were overriding feasibility problems, such as for Permanent Health Insurance. Nevertheless,

in other cases such as Job Shop Scheduling, a system was not implemented because the author was unable to convince the sponsor that there was a viable way to acquire the knowledge that would be both cost-effective and also would not take excessive expert time.

However, in many other cases, a successful expert system has been implemented, partially as a result of the perceived efficiency of the knowledge acquisition process. In this context it should be noted that "efficiency" can have two meanings and it is important to distinguish between them. There can be "efficiency" for the expert or "efficiency" for the knowledge engineer. Thus, in the first case, the acquisition process can be efficient for the knowledge engineer or researcher without being efficient for the expert.

For instance, many of the now "classical" methods of knowledge acquisition (eg repertory grid, card sorting) are quite efficient for the researcher but very inefficient for the expert. Both the techniques mentioned offer a very thorough and, for the researcher, comparatively straightforward mechanical route to eliciting the knowledge and, to a large extent, organising it into a convenient form concurrently with the elicitation process. (It can be argued that knowledge has not been "acquired" until the raw elicited knowledge has been understood and organised in some coherent way.)

Similarly, and as was the case in many early expert system implementations involving "dedicated experts", it can be quite efficient for the knowledge engineer to ask the expert to sit around, clarifying the occasional point, while a pilot or prototype expert system is refined in front of him. However, as with the "classical" methods above, such "on-line prototyping" is well beyond the time or tolerance of most experts. In fact, although we

respect these and other "literature" techniques, there has been no occasion when the practicalities of the real life situation have even allowed us to vaguely consider their use. We are convinced that techniques such as card sorting would be simply "laughed out of court" by most of the experts we have encountered.

Thus, rather than efficiency for ourselves as knowledge engineer, our first concern has always been efficiency from the expert's point of view. We have always tried to make the acquisition process the least time-consuming and easy for the expert. We have never attempted to use the more protracted techniques of the psychological laboratory and, by doing so, have managed to gain access to a range of expertise that we are sure would not have otherwise been available to us.

Nevertheless, as we have progressed, efficiency from the researcher's point of view has become increasingly important. Budgets for expert systems projects have become tighter. As a result, and perhaps by becoming a little more experienced in using quasi-psychological techniques in a way that is acceptable to the commercial expert, we have developed certain interviewing techniques for knowledge acquisition techniques that seem to offer reasonable efficiency for both researcher/knowledge engineer and expert.

Specific Techniques: Exploratory Knowledge Acquisition

Acquisition of the knowledge for any new task must inevitably pass through a series of stages. In the first instance, when the task is barely understood the knowledge acquisition will inevitably be of an "exploratory" nature and will be rather more unstructured than the later stages where

246

the knowledge engineer is "filling in" the knowledge on a more-or-less established structure. (This is the particular promise and merit of KADS, where once an interpretation model has been selected, the knowledge engineer can direct his acquisition techniques to the needs of the top-down process of populating the model.)

Whatever methods are chosen for the mid to later stages (and KADS seems a prime candidate), it is our assertion that the initial exploratory stage of knowledge acquisition is potentially the most inefficient and it is to this area that we have addressed most of our attention. There is good reason for this because, in our commercial work, it has been our ability to conduct rapid and efficient knowledge acquisition before the formal agreement to sponsor a project that has so often enabled us to make sufficient progress for us to prove the merit of the work to the sponsor.

In so many cases there has been a "Catch 22" situation where the reluctant sponsor has only agreed to a few hours of his precious expert's time for us to conduct our "experiments". Unless this time can be used to the maximum effectiveness then there will be no further opportunity to visit the expert to obtain more evidence on which to base a commercial proposal to proceed. (The problem is actually often worse than this because the initial knowledge acquisition must often also be used to "scope the knowledge" in order to provide a cost quotation for the building of an expert system. Underestimation errors at this stage can result in considerable financial penalties for the commercial implementer.)

Knowledge Acquisition – A Natural Ability?

A key personal belief underlying our knowledge acquisition techniques,

especially those used at the crucial exploratory stage, is that as humans we are naturally rather good at knowledge acquisition. Whether we are asking route directions or trying to understand someone's job at a party, the questions are easily and efficiently generated and, in the main, we get the information that interests us. It therefore seems unfortunate that many intending builders of expert systems are told that knowledge acquisition is difficult and a critical stumbling block in the expert systems and related areas.

Our personal feeling is that knowledge acquisition (ie primarily knowledge elicitation) is not too difficult. What is difficult is to make the process efficient and, perhaps more importantly, to follow it with a proper and efficient structuring of the acquired knowledge (eg such as offered by KADS).

The following techniques are amongst those used by the expert particularly at the exploratory stage of knowledge acquisition.

Specific Technique: Preparatory Briefing of Expert

The first technique is actually used in advance of the first formal meeting between researcher and expert. We have found it useful to avoid meeting an expert "cold" on the first knowledge acquisition session and, wherever possible, we have telephoned or, on one or two occasions, briefly meet the expert in advance.

The overt purpose of the ensuing conversation, rarely lasting more than fifteen minutes, is to give confidence to the expert. In most cases experts

are "volunteered" by their employers and are rather unsure what to expect. A friendly phone call will usually dispel any reservations and establish a basis of relaxed rapport for the initial session. We have discovered that experts, especially at the technician level (eg for the engineering tasks investigated), were often apparently concerned that "poor performance" at the knowledge acquisition sessions might affect their advancement. For this reason we have always given assurances that the results of the exercise will not be used for appraisals etc. Later, during knowledge acquisition itself the presence of the line manager will only increase such fears and it is better if he is tactfully excluded. A variation of this problem occurred in Pensions Management where, as explained in the Chapter 3 account of this task, there was considerable inter-expert disagreement. Before reconciling this problem by getting the sponsor to appoint a single authoritative expert, it was necessary to assure some of the experts, who were very experienced but junior to others, that there views would not trigger a dispute to their career detriment.

In the pre-meeting conversation, an explanation of the likely course of events has had a number of advantages. Apart from confirming "domestic" matters as time and place and reassuring the expert, such early contact has often have unexpected windfalls. In at least several instances experts have brought a pile of "prime examples" to the knowledge acquisition sessions.

For example, in one of the unreported expert system implementations following the Benchtop Electronic Diagnosis research the expert collected and prepared a series of "key example" defective circuit boards in the two weeks prior to our meeting. Similarly in Small Business Guidance the senior expert involved arrived with a collection of balance sheets with which to demonstrate particular key indicators of business health or otherwise. In

the case of Engineering Preventative Maintenance, and again as a result of investing a little time in prior contact, the expert spent several hours in private preparation and came to the first session with elaborate diagrams and notes. Such neatly typed "homework", complete in a binder, contained most of the knowledge of the domain. So, as a result of a few minutes preparatory conversation, a large part of our acquisition work was unexpectedly done for us by the expert.

Examples from our own research such as those quoted above show the same advantages as cited by Gammack and Young (1985) and Gammack (1987) who cite the value of asking the expert to prepare an introductory talk which, when delivered by the expert, is tape-recorded for subsequent analysis by the knowledge engineer. We believe that experts usually know the most important aspects of their knowledge and are often better at highlighting these than the "domain ignorant" researcher. In the early days of expert systems, some of the more successful systems have been constructed by experts themselves. Where this is not feasible, and the expert is sufficiently well-organised, self-preparation of "domain notes" is sometimes a powerful "second best" to self-construction of the actual system.

Specific Technique: "What do you do?"

This very simple and obvious opening gambit can often trigger the expert into a very full and competent description of the task which he undertakes. It is, after all, the question which he has been mentally preparing for in his anticipated encounter with the researcher. However the "What do you do?" question sometimes seems to be missed by workers anxious to rush into structured acquisition. For example, in the case of Commercial Business

Underwriting which, as we shall show, is largely procedural, the "What do you do?" question elicited the core "algorithm" underlying the expert's task. The expert spoke unprompted for about two hours after which it was merely necessary to "flesh out" the knowledge with specific instances of risk and their consequences. Again, although the "talk" was not specifically requested, there is a parallel to Gammack and Young (1985) and Gammack (1987) cited above.

In contrast the "What do you do?" question was almost worthless for Permanent Health Insurance and Eurobond Warrant Trading where it elicited little more than an elaboration of the expert's job title.

Specific Technique: Recognising Expert's Ability to Explain the Problem Domain

A particular goal of our knowledge acquisition has always been to acquire an early impression of the scope and arrangement of the domain from the expert's point of view. Almost without exception in the tasks investigated, this scope and layout has been different from that inferred from the personal preparation or sponsor's briefing. For instance, in Permanent Health Insurance the briefing indicated that the expertise was primarily in medical interpretation. In the event, it was largely in the area of assessing occupational consequences of disabling conditions.

As a result of this regularly found disparity between sponsor and expert, and particularly where the "What do you do?" question gives a shallow response, it is usually very worthwhile to ask the expert to recount a typical case. Often the best way to do this is to ask the expert to recount

the specific case he was dealing with immediately prior to the knowledge acquisition session. We have come to refer to this as the "yesterday" question.

The most recent or typical cases can be recounted in the richest detail. There is no reason to ignore this easily obtained part of the expert's knowledge by commencing with more obscure areas of expertise. In short, the "low-hanging fruit" should be collected first. In doing so the expert's confidence (and often his co-operation), is also usefully boosted by his ability to explain the more easily recalled material first.

The resulting torrent of material usually contains much problem solving "structural" material together with fluent details of specific cases and the knowledge relating to them. In Benchtop Electronic Diagnosis the "yesterday" question prompted the response "Another one of those funny little red ones". The peculiarities of faults arising from failures in a group of thirteen small red capacitors were to form the principal and most complex element of the acquired knowledge, and the "yesterday" question enabled most of this material to be acquired rapidly on the first day.

Sometimes a more structured form of knowledge acquisition actually prevents experts telling the knowledge engineer what he needs to know. Given the opportunity, most experts will cover the key elements of the domain first. If the structuring is left to the knowledge engineer then vital material may appear late or never. If later chided for his "omission" the expert will simply reply "Well, you didn't ask me". This has happened several times to the author and he has some concerns that, for instance in the later stages of a KADS-supported implementation, the knowledge engineer's questions may be so dedicated to the purpose of satisfying the interpretation model that

he may not recognise, or even may suppress, key knowledge that does not fit the demands of the model.

Specific Technique: Task Triggers and End-Points

A simple question that can usually be cleared up very rapidly is to ask the expert what triggers his activity. Most of the experts we have encountered work on discrete jobs. For instance for Telecom Repair Service the individual tasks are represented by particular faults on telephone lines; for Audit Work Planning individual tasks are represented by particular annual audits and for Life Underwriting each task is represented by a life to be underwritten.

Once the task trigger is established it is useful to ask the subsidiary questions on how the "input stack" is prioritised and what are the goals to be achieved from each job. Do such goals vary from job to job? In Pensions Management for instance there were at least two different sorts of "top-level" goals. An employee would either want an answer to a specific question ("Have I excluded my former wife from pension entitlement?) or would need a general overview of his likely benefit situation.

Similarly, it is useful to ask what was the end-point of each task. In Commercial Business Underwriting the task was completed when a competent quotation could be dispatched to the Branch Office. In the engineering maintenance and repair tasks completion was usually indicated by effective repair of the equipment. Apart from seeking the nature of the task end-point it is thus usually easy to determine the criteria that the expert, or his supervisor, use to determine when the task has been successfully addressed. Nevertheless, in cases such as Small Business Guidance the

253

end-point criteria has been quite vague.

Specific Technique: Easy, Difficult and Default Cases

As stated above, we strongly advocate allowing the expert "to tell his own story". Nevertheless, one of the most confusing situations that can occur in the early exploratory stages of knowledge acquisition is for an expert, perhaps as a result of his particular personality, to plunge into voluminous detail on obscure cases. This produces just the sort of detailed knowledge that is useful in the more structured later stages of acquisition. However, it confuses the exploratory stages when the researcher or knowledge engineer is still attempting to understand the gross structure of the knowledge. (Or, as explained above, trying to use his strictly limited time with the expert to assemble the best overall understanding of the domain in order to make a costed proposal for the construction of an expert system.) In the case of KADS, where an interpretation model has yet to be selected and there is no proper "host" for the knowledge, the researcher may be unable to properly deal with the detailed outpourings at this stage. Since experts seem reluctant to repeat things that they feel they have already explained in great detail, there is a consequential danger that this ill-timed knowledge may be lost or may be difficult to analyze during the subsequent knowledge acquisition sessions.

In such cases we have found it useful to use specific questions which act as probes to focus the knowledge acquisition in a way similar to that described in Burton & Shadbolt (1987). For instance, a useful strategy is to interrupt the expert and say, "Tell me, what would happen on an easy day?". For instance, when we used this method in the investigation of Small

Business Guidance we received the response, "Well, a nice cashflow problem would be a good start". When this and other straightforward cases had been explored and the structure of the domain understood we were more able to cope with the expert's digressions into some of the more complex areas of his knowledge.

In contrast, some experts try to be helpful by never dealing with anything other than simple cases. One of the subsidiary experts supporting Audit Work Planning was in this category and perhaps this was due to the accountant's instinct that nobody outside their profession can appreciate anything except the most simple examples. A similar situation occurred in Foreign Exchange Exposure and, as with accountants, it is best to ask a "difficult case" question such as "What is the worst type of case where you have to offer currency exchange rate exposure protection?" The response "Turkish lire, anything over 18 months....." and immediately moved the expert into more tortuous and productive areas of his knowledge.

Related to the use of easy and difficult cases, we have frequently asked experts to describe their response to what we refer to as "default cases". Two categories of default cases are proposed, common cases and the empty case.

There is an adage in the expert systems community that "ten per cent of the knowledge can cope with ninety per cent of the cases". This may not be strictly true but it does seem that most tasks involve a high proportion of regular cases, which most staff can tackle, plus a much lower proportion of more obscure cases which are dealt with by the more experienced staff.

Although they represent a large segment of the domain knowledge many of

the obscure cases have only been encountered two or three times previously. In essence this less common knowledge is the knowledge that distinguishes the expert from his colleagues. Nevertheless, to provide a framework for the researcher's understanding and, probably, the final system it is useful to capture the ten percent of mundane knowledge first. The highlighting of common and empty cases serves this purpose.

Common case capture is straightforward. The expert is asked to list the three, five or ten most common cases which he encounters in his work. Then, just as with the use of "yesterday" and "easy day" cases described above, he is asked to describe his activity if one or more of such cases occurred.

The description with such self-declared cases has always been fluent and provides a very rapid and easy method of eliciting much of the bulk of the routine core knowledge of the domain. Moreover, when the expert is in such confident ground, it is far easier to acquire peripheral knowledge on tools, documents and other props without the risk of disturbing his primary dialogue.

The empty case tactic is a little more contrived but results in a very useful input for when the knowledge is later structured. The expert is asked to describe his activity if an "empty case" was put in his in-tray or workpile. In engineering terms an empty case is, for instance, a fault-free unit in a pile of faulty units for repair. In commercial terms an empty case could be a straightforward healthy risk-free applicant in the midst of doubtful cases requiring life or permanent health insurance.

The value of the empty case is that it causes the expert to go through his

default search routine of which, in many cases, the entire knowledge is an elaboration. Obtaining this key structure directly and easily can be particularly valuable. It was used widely in a number of forms for Telecom Repair Service, Benchtop Electronic Diagnosis and Small Business Guidance.

For instance in Telecom Repair Service the RCO was asked "Suppose a customer reported a line that was actually in perfect working order?". In Benchtop Electronic Diagnosis the question was "Imagine that your supervisor put a non-faulty power unit on your bench, what would you do?". In Small Business Guidance the business guidance consultant was asked "What would happen if you were asked to provide guidance to a small firm that eventually proved to be healthy and prosperous?

Specific Technique: Counterexamples and Hypothetical Scenarios

These two well known techniques have been widely used both at the exploratory and more structured later stages of knowledge acquisition. At such later stages they will usually be used to follow up knowledge analysis where it has been discovered that an apparently feasible problem situation has not been discussed with the expert. For instance in Telecom Repair Service the expert mentioned a range of categories where earth and battery faults of different degrees of severity (ie different impedances to full battery or full earth) could affect external cables. On later analyzing the acquired knowledge it was realised that only a fraction of the possible permutations had been covered in detail by the expert and therefore these were used as the basis of further questions to the expert.

In the exploratory stage of knowledge acquisition we have usually used

counterexamples on an immediate "on the fly" basis where it is realised that the expert is only covering one out of two or more options. This situation was particularly prevalent in the knowledge acquisition for Small Business Guidance where there were always many alternatives to the specific case details discussed by the experts.

The use of hypothetical scenarios is somewhat different from the use of counterexamples. In the main only one element of the example is changed in a counterexample situation. In a hypothetical scenario a completely new set of problem elements is usually presented. This latter technique has proved most valuable to assist experts to "move laterally" to other areas of their expertise. In Burton & Shadbolt (1987) such examples were usually prepared by the knowledge engineer prior to the session with the expert; in our work we have used the technique on a purely opportunistic basis with example scenarios developed by the experimenter during the course of a particular session.

Specific Technique: User Knowledge Acquisition

The purpose of an expert system is to offer an expert's knowledge to a "user". In different expert systems, and sometimes for the same system, "users" may range from novices, with little knowledge of the expert's domain, to other experienced practitioners who just happen to lack the specific experience of the "donor" expert. (An informal classification of "experts" and "users" is discussed in Chapter 6 to demonstrate the different factors that these different groups dictate for the success or failure of an expert system.)

258

However, in all cases there is an implicit requirement to make good some difference between the expert's and users' knowledge of the domain. It is still often naively assumed that experts know, first, what users need to know and, second, how it should be delivered to them. We discuss these two issues, of "content" and "style", below.

Users are the real "customers" for expert systems. Expert systems exist primarily or solely to satisfy users' needs. It is therefore strange that so much expert system building methodology concentrates on acquisition of the expert's knowledge without any proper assessment of the nature of the users' need for that, or other, knowledge. (Nevertheless, there are exceptions such as the stress on determining users' needs in the work of Mittal and Dym, 1985.) The preoccupation with expert knowledge at the expense of user requirements is a key distinguishing feature between expert systems methods and the methods employed for conventional software development. In the latter it is the business, including user, needs that drive the process rather than the availability of "knowledge".

Perhaps the emphasis on "expert push" rather than "user pull" is one of the contributory factors to the poor take up of expert systems technology. It is no use making the most elegant and clever expert system manifestation of an expert's knowledge if there is no business need for the knowledge to be so delivered.

Thus we suggest it is a prime responsibility of the expert system builder, perhaps even before he or she starts to seriously acquire expert knowledge, to make a proper assessment of the user space that such knowledge is intended to fill. Thus the "content" of the knowledge which an expert system delivers should match the users' actual need for knowledge rather

than, as is often the case, just being a replication of the total knowledge acquired from the expert.

Further, an expert is recognised by his task competence, not as a result of his competence to articulate his knowledge in a convenient form to a less experienced person. Just because an expert is the custodian of useful knowledge there is no guarantee that he can deliver such knowledge in just the way that will most ably assist the novice. As with so many other teacher/pupil situations, the delivery style of the expert may not match the required input or cognitive style of the novice. We do not claim to have made a proper investigation to confirm the desirability of matching the expert's style and user's requirements. Intuitively we feel that such match is desirable but acknowledge that this an area that requires proper investigation. This research has regretfully not afforded the opportunity for thoroughness in this respect.

It is our contention that experts can usually be relied on to offer the content of the knowledge (subject to the "user need" comments above) and that users can be largely relied upon to advise upon the style in which they require knowledge to be delivered to them. For instance, in Pensions Management, users were very clear on their requirements even if these requirements were often articulated in terms of what was wrong with the style and knowledge delivered by Pensions Officer experts.

In view of the content and style issues just raised we see a very definite need to undertake some form of "user knowledge acquisition". (More properly we are acquiring the user's need for knowledge and the style in which it should be delivered to them.) As our commercial expert system work has progressed, and become more successful, we have confirmed this

need and the practical advantages of user knowledge acquisition. A system that answers a well-determined and pre-specified need in the user community is much more likely to be accepted than one that just offers a clever but, at worst, gratuitous delivery of the expert's knowledge.

To emphasise this point it is now common for us in our commercial work to spend up to 40% of our knowledge acquisition time allocation on user knowledge acquisition.

Usually it is fairly easy to identify the expert or experts who will contribute to the content of an expert system. However, in a one to many situation, the number of potential users will greatly exceed the number of experts. For this reason there must be selection of a representative set of potential users.

For instance in Pensions Management it soon became obvious that both Pensions Officers (local "experts" needing the support of the key Pension Scheme experts) and personnel/pensions clerks (clerical staff with only rudimentary pensions knowledge) would be required to use the eventual expert system. In these cases the user groups differ in their level of pre-existing knowledge and familiarity with domain language (ie "buzzwords") and concepts. Distinct differences in the level of knowledge of different expert systems user groups (often corresponding to different staff grades) are now increasingly recognised. The eventual user panel for Pensions Management therefore consisted of three Pensions Officers (of differing experience), two pensions clerks (of different seniority) and a personnel clerk who represented occasional "low-level" potential users of the system.

A usual minimum figure of three potential users is based on a pragmatic compromise between the need to give some sort of representation to different user styles and the need to maintain the cost of user knowledge acquisition.

User knowledge acquisition usually lags behind the acquisition of expert knowledge and is best conducted in parallel to the later analytical stages of expert knowledge acquisition. In KADS the users' needs are partially covered by the Model of Cooperation although we suspect that this may usually be based on the expert's view of what users need rather than their own expressed requirements.

In all cases an assumption is made that users can solve some of the easy cases but must consult or refer the problem entirely to more experienced colleagues (ie the expert or his immediate colleagues) to resolve all or some of the difficult cases. There is thus a key requirement to note the points of access to the senior expert, the type of advice sought (ie the user's need) and the form of the response from the expert.

By conducting such analysis in the "human expert" situation, prior to the introduction of an expert system, it is possible to classify users into two approximate categories, "consultation users" and "refer to expert users". In the cases of "consultation users" the type of advice sought may be a specific clue or instruction on how to proceed or complete guidance on how to solve the problem to its end.

Alternatively, and in the majority of cases, the "refer to expert users" require that difficult cases are completely referred to human experts who take over the cases and complete them through to completion. In such

cases the users are often only left to close the paperwork, in clerical tasks such as Life Underwriting, or dispatch repaired equipment, as in the case of Benchtop Electronic Diagnosis.

Thus, for clarification, some users of expert systems will have never even attempted to undertake the tasks for which the expert system will support them. Usually they just refer such "difficult" cases to a more senior colleague. However, when supported by the expert system, they will have to carry the case through to completion and, since they will be "treading new ground", this imposes a very particular requirement on the expert system.

Specific Technique: Prototyping to Assess User Requirements and Acceptance

Regretfully the user population is invariably unable to define what (and how) it needs to know. Expert systems are still relatively new and users have little experience on which to base their requirements. This is in marked contrast to conventional computer systems where "non-computing" users have become relatively sophisticated in their requirements. Users of systems ranging from simple spreadsheets to major interactive business systems now often have sufficient experience of existing comparable systems to quite clearly define their needs. However, for the clarification of expert system user requirements, it is usually necessary to resort to some form of prototyping.

Structural prototyping is used to define the best overall plan of the delivered system. For instance, in the case of Small Business Guidance it soon became obvious that potential users, who were usually owner-managers

of small businesses, would not tolerate a long exhaustive preliminary fact-finding dialogue on their business.

Instead they wished to provide basic details and then rapidly enter an interactive diagnostic on specific business problems. The end result was the same because the interactive diagnostic actually required the user to provide the same core factual information but, because it was interspersed with user feedback, the users would tolerate its request.

However, consensus on user requirements is not always easy to reach where the user population is at the more experienced end of the spectrum. In tasks such as Audit Work Planning, Multinational Tax Planning and Pensions Management we have had considerable debate at this stage where "near-experts" (or experts not selected for knowledge acquisition) have argued with the content of the knowledge provided by the nominated experts. In these cases the issue really is one of corporate choice of the knowledge to be offered by the system. It may be that equally well-esteemed auditors or tax practitioners or pensions experts can actually derive and offer their advice in completely different ways. An expert system can usually adopt only a single style of operation and, in such cases, there must be a clear means of resolving these issues at a senior level with the system sponsor.

User interface prototyping is arguably the most important part of user knowledge acquisition. The user interface is the only part of an expert system (or any other computer program) seen by users and the acceptability of the overall system will often be judged on user interface issues. Smith & Mosier (1986) provides a useful and comprehensive set of guidelines for user interface design.

Prototyping is most strongly evident here as it is both the quickest and most direct way of determining user needs. User interface prototyping can be undertaken in isolation from the rest of the system design and implementation. For instance in Audit Work Planning, which was eventually implemented using an in-house proprietary shell, the user interface was initially prototyped using just a fast screen painter tool. Such methods follow the general ideas expounded in Bricklin (1986).

By chaining together a series of totally non-interactive screens it was possible to fully simulate a particular audit planning sequence merely by hitting the "Return" key. This simple measure, well in advance of the finalisation or implementation of the knowledge base, enabled potential users to offer many comments on the required appearance of the system. Similarly for Multinational Tax Planning the appearance of the interactive balance sheets was prototyped well before the complex task of actually implementing them was undertaken.

Another important user requirement concerns the form of any reports produced by the system. In the case of Life Underwriting the system produced screen reports for other users of the mainframe hosted system. Such reports had to be precisely in the same format as reports generated by a human underwriter. Even more stringent was the user requirement for Audit Work Planning where the actual audit plans had to be on identical paper and laser printed to be similar to the printed "manual" audit planning forms used previously. It was only by this route, and by allowing for the system generated audit plans to be hand signed, that audit staff would accept them as authoritative.

The reports for Business Performance Assessment had to be of a similar

quality to those prepared manually since they were to be submitted for Ministerial comment. As detailed in Chapter 3 they were sufficiently good to be confused with the manually prepared reports.

Specific Techniques: Knowledge Capture

In our initial commercial work we adopted the procedure of collecting, transcribing and analyzing complete transcripts from knowledge acquisition interviews. This, of course, had been the preferred method of most cognitive psychologists and seemed "correct" for our investigations.

Nevertheless, as our work expanded, we soon realised that the process of preparing and analyzing complete transcripts was over-laborious and inefficient for large segments of the investigation of the various tasks. Accordingly, we adopted a more pragmatic regime for the capture and subsequent processing of elicited knowledge. Such subsequent processing invariably did not involve the production of either protocol or interview transcripts and, in main, systems were built on the basis of knowledge directly "transcribed" from audio or video tapes.

In all cases it is now our rule to make a complete audio record of all knowledge acquisition sessions. Able colleagues have sometimes abandoned this practice and, for the skilful taker of legible notes, there is still some justification for making direct notes of an interview.

However, in our case, we have always found that the actual business of conducting the interview to be sufficiently difficult so as to leave little capacity for making a competent record. Also, where colleagues have

prepared direct notes, knowledge acquisition has been rather more protracted and, in some cases, errors and omissions have become noticeable.

Video recording has been used to supplement audio recording on a number of cases. It is always rather more intrusive than audio recording and we now only employ it for tasks where we know that it will be positively useful. In particular, video recording was used during parts of the investigations of Benchtop Electronic Diagnosis, Audit Work Planning, Multinational Tax Planning and Eurobond Warrant Trading. Our usage is similar to that of Dym (1987) who recommends that certain tasks are best investigated by asking the expert to perform the task, usually in front of a video camera, rather than describe it.

Video recording is particularly useful for tasks where the experts use what we refer to as "fingerthinking". In such tasks an expert will usually explain his reasoning by tracing his finger over a diagram, chart or "paper spreadsheet". Thus, for instance, the electronic diagnosticians that we have encountered invariably explain their circuit testing and analysis using a circuit diagram. "If this component X is defective then an abnormal current will flow in this direction and can be detected at this point Y" is typical of the remark that will be accompanied by a tracing of the current path by a moving finger.

Similarly, financial "balance sheet thinkers" (eg tax and audit specialists) will usually want to illustrate comparisons and ratios by direct pointing to familiar critical areas of the balance sheet. We suggest that such props, such as circuit diagrams and balance sheets, are key organising devices for the experts and it is merely a secondary but useful characteristic that they are used for explanation. Their main purpose is as "problem solving props",

as discussed above, to directly assist the experts in their problem solving.

It has been found that video recording is best undertaken with the camera in one of two positions, "top-down" or "over-the-shoulder". In both cases the camera is directed at the work area and, in particular, never is directed into the faces of the expert. The top-down position is most suitable for viewing circuit diagrams or sheets of figures etc which are placed in the "blotter area" of the desk or bench at which the subject works. In such circumstances it is usually possible to mount the camera tripod at the back of the expert's desk or bench with the camera pointing more-or-less vertically downward. The researcher can then sit to one side of the expert to conduct the interview with the expert who can "fingerthink" across the documents immediately in front of him.

The over-the-shoulder camera position is used in the workbench situation where the expert is pointing to instruments, and the unit under test, over a "benchtop panorama". In this case the researcher can actually stand in front and to one side of the camera and conduct the interview at the same time as panning the camera from one area of the bench to another. A video camera with automatic focus and a forward facing pan handle are required for this form of operation. Although quite difficult in terms of manipulation the visually acquired detailed knowledge can be considerable in appropriate cases.

Specific Technique: Knowledge Acquisition by Review

Knowledge acquisition for Eurobond Warrant Trading was undertaken in a form that the author refers to as knowledge acquisition "by review". This

form of knowledge acquisition is similar to that used for acquiring medical decision making knowledge as, for example, Elstein, Schulman & Sprafka (1978).

There are many similarities between the normal working environments of the experts for Telecom Repair Service and Eurobond Warrant Trading. Both experts work with the aid of information technology (computers for both tasks, test desks for Telecom Repair Service, information feeds (eg Reuters) for Eurobond Warrant Trading) and conduct much of their business by telephone.

However, although the Telecom Repair Service expert was under pressure to diagnose faults, the Eurobond Warrant Trading expert was in a real-time situation where large sums of money were involved. As a result, although we could undertake "online" knowledge acquisition in the working environment for Telecom Repair Service, this was not tolerable for Eurobond Warrant Trading. (The same situation applies for clinical diagnosticians who are often involved in real-time life-or-death decision making and will not tolerate interruptions.) Nevertheless, as with Telecom Repair Service, the use of case protocols and case-related examples appeared to offer a very productive method of knowledge acquisition. In fact knowledge acquisition for Eurobond Warrant Trading without access to "live cases" did not seem to offer very good prospects for success.

Consequently it was necessary to devise a method of knowledge acquisition that was sensibly non-intrusive in the working environment yet allowed the experimenter access to the very rich case material available. Knowledge acquisition "by review" was the successful method devised. The core idea of knowledge acquisition by review is to undertake a "passive video

capture" of as much as possible of the expert's working day and then use this as a "prompt" for off-line knowledge acquisition in the evening. In the case of the Eurobond Warrant Trading investigation, which is the sole instance for which we have used this method, a video camera was set up "over the shoulder" of the Eurobond Trader's position before the commencement of the Trader's working day. A microphone was placed on the Trader's desk where it could pick up his comments and his speech on the telephone. It was not connected into the telephone circuit since this would have involved recording a counterparty's speech without his knowledge or permission.

Apart from the occasional civility and comment, the experimenter did not intrude on the Trader's activity but merely ensured that the camera was positioned so as to pick up as much detail from the screens of the information feeds and the PCs running spreadsheets and other decision aids to assist the Trader. Both morning and afternoon sessions were recorded on two consecutive days.

Following completion of the afternoon sessions and a short break, the experimenter conducted a knowledge acquisition session with the Trader expert. These sessions, which lasted several hours, were held in a comfortable conference room with refreshments away from the normal trading environment. A large TV monitor was set up to view the video-recording of the same day's session. An audio tape recorder was also set up to capture the review session speech of both expert and experimenter together with that of the video segments. The video tape was played in segments and was "interrupted" (as in the "online" case of Telecom Repair Service) both between natural trading sequences and also where the experimenter needed to ask questions. Such questions and answers were, of course, fully

audio recorded

Typical examples of the questions and responses prompted by the ongoing
video sequence are:

1. E: Why did you raise your price at that point?

    S: I had had several separate price enquiries from Japanese banks
            during the morning and thought something must be "on the
            move".

2. E: Why did you ask him (a counterparty) for a price on that
            warrant?

    S: I have not traded that one for nearly a fortnight and I feel my
            own price may be to high. Also I was worried that I then had
            too many of those warrants and ought to get rid of them.

    E: And why are you worried about that warrant?

    S: Its price always runs parallel to XYZ warrant and that was
            being offered at a much lower price this morning. So I should
            lessen my exposure by decreasing my holding.

3. E: What were you thinking then?

    S: Well, I was losing money rapidly so I decided that the safest
            thing was to do nothing and hope that the market would
            recover later in the afternoon - which it did!

Questions such as those above provide powerful insights into the expert's beliefs and reasoning. Such insights could not have been directly acquired during the working day. Provided the review session is held shortly after the recorded session then the expert can remember the concerns and hypotheses that were actually current at the time the video sequence was recorded. As a result, coupled with the ability to pass over uninteresting sequences or "stop the day" at points of interest, knowledge acquisition by review provides an equal if not more powerful method than protocol collection or interview in the working environment.

Although Eurobond Warrant Trading did not result in an expert systems implementation, the experience gained in knowledge acquisition by review was very worthwhile. The technique is easy and not stressful or intrusive to the expert. The video record is a powerful prompt to the expert and enables him or her to provide a very rich source of knowledge with minimum extra time input to the investigation.

"Perceptual Knowledge"

The previous section dealt with the acquisition of knowledge for Eurobond Warrant Trading. This method was successful but we encountered occasional instances of what appeared to be "perceptual knowledge". Our use of the term "perceptual knowledge" seems to be similar to the knowledge of chess players who "recognise" board configurations as reported in Chase & Simon (1973). "Perceptual knowledge" may also be an example of what Dreyfus and Dreyfus (1986) call "arrational knowledge". They claim the highest of five stages of skill acquisition is "Expertise" and that, rather than being "calculative problem solving", it "involves critically reflecting on one's

272

intuitions". Dreyfus and Dreyfus cite references to expert X-ray interpretation, chess masters and airline pilots in support of their arguments in favour of arrational knowledge. Theirs is a controversial area and we would hesitate to engage it. However, we cannot ignore what we term "perceptual knowledge" and, apart from Eurobond Warrant Trading, we believe that it is actually more common in other trading environments such as Foreign Exchange ("FX") Trading.

This section is prepared on the rather limited investigation of Eurobond Warrant Trading but also builds upon considerable discussions with colleagues concerned with the difficulties of acquiring FX and similar knowledge.

FX Traders have great difficulty in articulating their knowledge and, if asked to justify an action, put it down to "experience" or "gut feel" etc. It is our suggestion that the knowledge of such experts, although originally available for introspection and acquisition, has actually become "compiled" so as to be a form of "perceptual knowledge". Thus traders "recognise" a situation on their Reuters FX screens and act upon that "recognition".

Although knowledge acquisition from FX Traders and similar experts is very difficult, the potential resultant system rewards are very high. We suggest that techniques such as inductive knowledge acquisition might be used in such cases where knowledge cannot be acquired directly. Inductive knowledge acquisition, and routes to increasing its utility, is discussed in the next section.

Specific Techniques: Inductive Knowledge Acquisition

Expert systems induction tools have a long, but not very successful, history of application to classification problems. Much expert decision making can be seen as a process of classification and, as a result, the usage of induction is often proposed as a quick and effective means of knowledge acquisition based on the exploitation of existing decision data. The seminal work of Michalski and Chilausky (1980) on soyabean disease classification is the basis for the enthusiasm of many inductionists.

In conjunction with his commercial work, the author has had the opportunity to make some useful investigations into the utility of induction tools and methods and, since they apply more generally, the results are reported here.

Induction tools are claimed to have some of the following benefits for expert systems construction:

1) They enable existing knowledge engineers to build successful expert systems faster than by "manual" knowledge acquisition.

2) They enable less skilled knowledge engineers (or existing systems analysts etc) to build successful expert systems to the standard of skilled knowledge engineers.

3) They enable existing records of expert decision making (ie classification data) to be exploited to build an expert system without the co-operation of human experts.

4) They quickly intrigue and capture the imagination of commercial sponsors and experts who would otherwise not be interested in expert systems.

5) They provide a quick and effective prototyping medium to capture the core knowledge for systems to be finally completed by conventional expert systems methods.

Currently, most research and commercial applications of induction methods and tools seem to employ the usage referred to in (3) above, ie the exploitation of existing data on expert decision making. However, many potential (and often very enthusiastic) users of induction tools discover that, in reality, it is very difficult to find good examples of suitable data to be usefully analyzed. All induction tools come with interesting historical data sets (usually including the classic soya beans example from Michalski and Chilauski (1980) to demonstrate the utility of the tools. Nevertheless, in practice, users actually find few worthwhile examples of useful data in their own companies on which to use the tools.

As a result, without the ready availability of good examples of tabular data, the vendors of technically excellent induction products have faced commercial failure in selling their products.

Nevertheless, we suggest that the usage anticipated in (3) above represents a very large and valuable potential usage of induction technology. Such usage is actually feasible provided the definition of "existing records of expert decision making" can be modified to take in data that is not in a neat tabular form. This matter is considered in what follows.

Knowledge can be acquired directly from experts or induced from the accumulated results of their work. It was for the latter purpose that induction tools were originally intended. The developers of such tools have seemingly expected that the expert's decisions would be neatly recorded in a tabular form or could be easily transcribed from written records to a spreadsheet or database by a simple clerical process.

However, most expert decisions are not recorded in such a neat form. In many commercial tasks case data is recorded or accumulated (eg as secondary documents) in paper files. Such case papers are then examined by the expert to form the basis of the decision. Once the decision (ie categorisation) is made the result is also recorded in the case file. Some typical real-life commercial examples (including examples from the tasks reported in Chapter 3) are:

1) Trade credit assessment

2) Staff recruitment and appraisal assessment

3) Business performance and loan appraisal

4) Research grant applications and awards

5) Commercial Business Underwriting

6) Permanent Health Insurance

7) Foreign Exchange Exposure

In such cases it is not possible to simply encode all of the contents of all of the files and induce rules from the resulting mass of data. The key difficulty is that the expert may be very selective about the use of case data and may use low-level "data interpretations" that are not immediately obvious. These points may be made clear by the following real examples:

1) "Is the credit application form filled in neatly, untidily or with average care?"

This example is taken from the author's commercial work on Trade Credit assessment. Once the attribute and allowable values were identified by the expert it was easy to codify the "neatness" attribute on all the application forms into one of three categories ("neat", "untidy" and "average"). However it was not possible to identify this attribute, or its range of permissible values, without the advice of the expert. (In passing it is worthy of note that the neatness of the handwriting was a key indicator used by the credit control expert. "If he fills in the form carefully then he will be careful with the accounts". To the author's considerable surprise, more "scientific" and "statistical" methods of assessment took second place to the handwriting test.)

2) "Does the life insurance applicant play ice hockey, ski or fly gliders?"

This example is taken from Permanent Health Insurance where an underwriter looked for the answers to these simple questions amidst a mass of free format text on an application form. It was impossible to encode all the possible permutations of details on such a form but, once the critical data item was identified, it was a routine task to codify the existing mass of application forms.

If we accept that much real life data is in the above form we must conclude that induction must be preceded by a process of what we term "attribute acquisition". The need for such attribute acquisition, for which the author has undertaken some experimental work, is explained as follows.

277

Tidy tabular decision data, which is anticipated by induction tools but which hardly exists in real life, has the attributes and decision classifications already identified.

Real life decision data (which is abundant in real life) requires the critical attributes (and their allowable values) to be identified or highlighted by the expert before induction methods and tools can be applied. Attribute acquisition is thus a key technique to enable the useful exploitation of inductive knowledge acquisition.

Where attributes are acquired as a means of codifying existing real life data (or as part of iterative knowledge acquisition as described below), the techniques used will differ from those used in conventional knowledge acquisition. In the main attribute acquisition is a much "shallower" activity than knowledge acquisition. The primary challenge is to identify attributes but prevent the process from escalating into full knowledge acquisition. If such escalation occurs there is a considerable temptation to "manually build" the resultant expert system. The process also becomes heavy on expert and knowledge engineer time and commitment. It is, of course, to avoid such manpower effort that induction is often attempted.

We have undertaken experimentation with techniques for attribute acquisition but, although promising, the results are not sufficiently advanced for inclusion in this thesis.

4.4   PROBLEM SOLVING PROPS

Introduction

Professional and technical staff in commercial organisations solve problems on a routine basis.   Most of the "experts" we have observed, both as reported here and in our wider work, solve their everyday work problems without undue stress or painful "exercising of judgement".

Yet their activity is not trivial and cannot be undertaken by laymen.   They are experts, at least in the eyes of their colleagues.   However part of their problem solving adroitness results not so much from their own knowledge but from the problem solving "props", or tools, that they use in their work. (We shall use the term "props" rather than "tools" since this latter word is used more commonly in this thesis to refer to expert systems tools.)   These props are either provided by their employer (as part of the institutionalised problem solving framework) or are evolved and used on a "custom and practice" basis by the experts themselves.

Most problem solving props serve the same purpose.   They assist the expert to readily cope with the majority of problems in a well-ordered, thorough and effective way.   If the problem is outside the scope of the prop, or the expert's personal experience, then the expert must either work from first principles or seek the help of a more experienced expert.

Props can lead us to understand how experts solve problems and are thus a powerful aid to knowledge acquisition.   They also provide evidence for weaknesses in human problem solving.   By "weaknesses" we refer to those

general limitations of human cognition rather than the specific failings of less able individual humans. We suggest that the props are often used as methods of "patching" such weaknesses or for allowing an expert to solve problems beyond his or her normal capacity.

A common non-occupational example of the latter is the use of pen, paper and a standard format to allow us to undertake multiplication tasks of greater complexity than those possible by "mental arithmetic". The paper provides us with an "external memory" to hold the partial results which are beyond the reliable capacity of our short term memory (STM).

The study of problem solving props is also valuable in that they provide insights into the type of support which could be provided by more sophisticated problem solving aids, such as decision support systems (DSSs) etc. Thus if we can analyze the role and function of existing props then we may be able to see what roles and functions are feasible for "intelligent props".

The following are some of the more common props encountered in the investigations of occupational problem solving reviewed in Chapter 3.

Checklists

Probably the most common, and simplest, employer provided tool is the checklist. This is all-pervasive in routine tasks not classifiable as problem solving (eg packing lists, routine maintenance checks etc). Nevertheless it is also widely used to support straightforward problem solving. "Check battery; check fuses; etc" is one of the most common beginnings for

electrical troubleshooting guides for cars, radios etc. Most such checklists, when designed for end-user non-experts, finish with a sentence such as "If failure persists then seek professional advice...". Checklists for industrial/commercial use have a similar referral to a senior expert or specialist. "For risks in excess of £100,000 refer to Mr Jones for approval/advice" is typical of the referral to a named senior underwriting expert.

Checklists were very much used by the junior "clerical experts" in Pensions Management. When an employee asked for an official statement of his future pension entitlement (or when a retiring employee was transferred to pensioner status) the clerk would use an employer-provided checklist to ensure that the employee had satisfied all the detailed conditions of pension entitlement. This was essential since a mistake could offer the promise, or actually instantiate, a pension not allowable under the pension scheme rules. Management would not have been satisfied with a "mental checklist" and staff had to tick the specific items on a form as they were checked. Naturally, this checklist was one of the most essential, and easy, elements to incorporate into the delivered expert system.

"Virtual Props"

In some cases the "props" do not physically exist before the commencement of the solution to each problem. They are created over and over again, like a form of mental scaffolding, as each problem is solved. An example of this type of "virtual prop" is the structure of the quotation Commercial Business Underwriting. The structure of the "setting out" of the quotation calculation, even though it is completed on blank paper every time, is the

prop which helps discipline the solution to the problem. It is exactly analogous to the standard format of the schoolboy's geometric proof which, although it serves a multitude of different proofs, is the essential route to a disciplined solution.

Paper Based Props

Checklists, already mentioned, are often paper based. However checklists are the simplest of a whole range of paper based props and rarely demand more than just a series of ticks, either mental or by pen. Most paper based props support and structure more specific aspects of the expert's problem solving.

For instance, in Telecom Repair Service the "docket" used to record the customers report and, subsequently, the fault diagnosis, forces the technician to classify the report into one of a series of predetermined categories. Although the vague "OOO" (Out of Order) is an allowable classification, it is last on the list and its use is definitely not encouraged. This "classification" form of prop very obviously serves the purpose of disciplining the problem description. It provides the diagnostician with a complete range of descriptions without which it is very likely he would use either a more restricted list or move straight into diagnosis without first clarifying the problem description.

The docket is the simplest type of paper based prop that can be classified as a "form" rather than a checklist. In Audit Work Planning a much more complex type of form was used. It served both checklist and classification purposes. For instance, and to satisfy the legal requirements of a statutory

audit, the form cannot be properly completed until a number of categories of information about the audited firm have been properly recorded. This is essentially a checklist function. In terms of categorisation function, it causes the audit manger to classify the firm's trading activity, type of computer support etc into a set of pre-defined categories.

Additionally, and beyond a straightforward checklist function, the form used in Audit Work Planning also contains some elements of an algorithm. Early questions on the form incorporate a sort of "GOTO" indicating other areas of the form to selectively complete. (UK passport application forms are similarly designed.) Such considerations almost invariably involve the collecting of further information. In this respect the problem solving prop does not lead the manager to the details of the audit planning "solution" (which remains his task based on his experience or, since our implementation, the expert system) but force the manager to provide himself with the same sort of information that the form designer knows will be necessary for him to competently plan the audit. We classify this as a Global Data Acquisition component of problem solving. In the implemented expert system such function was more-or-less directly replicated as a straightforward interactive questionnaire module.

Graphical Props

Job Shop Scheduling and Multinational Tax Planning demonstrate the use of different forms of "graphical props". In Job Shop Scheduling schedulers almost invariably use a form of Gantt chart in order to structure their thinking and "see" the interrelationship between competing orders etc.

In Multinational Tax Planning the corporate organisation chart is used (or often partially sketched for each problem) to allow the tax specialist to take proper account of the various tax dependencies.

In both these examples, Job Shop Scheduling and Multinational Tax Planning, it seems very likely that the problem solver could not handle the problem if using a purely text-based prop. A prop showing the "spatial relationship" of events (in a schedule) or entities (for subsidiary companies) seems essential. We have taken due account of this in the design of the graphical interface for our Multinational Tax Planning expert system and, for exploratory work in the job shop scheduling area, fully realise the importance of the Gantt chart as part of the system's user interface.

Decision Support Systems

Decision Support Systems (DSSs) represent a very special form of computer based prop used by experts. Their primary advantage is that they can be caused to display the analyses of large amounts of data in a form easily comprehensible to the expert. In addition they can allow the expert to perform "What-if" experiments on the data.

Except for use by the expert for Eurobond Warrant Trading no examples of DSSs were encountered in any of the tasks investigated in this thesis. Nevertheless the behaviour of experts using DSS systems is almost uniquely traceable (ie by means of a trace file) and could provide a very convenient asset in knowledge acquisition. In theory such usage traces would be very able support "knowledge acquisition by review" discussed earlier in this chapter. In the future, where experts are supported by hypertext systems,

the usage record could also provide a good record of the information sought by the user to support his problem solving and, coupled with "review" knowledge acquisition, could be invaluable for expert systems building. For instance, the risk manuals and specific application details for Commercial Business Underwriting are all paper based. If an hypertext system had been used, as is likely in the future, then it would have been relatively easy to infer the expertise used largely on the basis of "re-playing" the expert's usage record.

Technical Props as a Help Implementation

In Telecom Repair Service the "tool" was complex and embraced the entire Test Desk. In Benchtop Electronic Diagnosis the tools were primarily screwdrivers, soldering irons etc and test instruments such as oscilloscopes etc. In Commercial Business Underwriting the props were the rates, manuals and special conditions manuals. At the exploratory and intermediate stages of knowledge acquisition discussion of the "props and tools" can be a very rapid way of eliciting considerable knowledge about the task. In the later stages of knowledge acquisition it has usually been necessary to return to a more complete description of method of use of the "props and tools" for the purposes of the constructing the delivered system's "Help" information. (For instance, in Benchtop Electronic Diagnosis, "Help" screens were offered to provide details of the settings and use of all the testbench equipment.)

Problem Solving Props - Summary·

Many occupational tasks demand the use of tools. In the manual tasks these might be soldering irons, screwdrivers etc. (Their study can be invaluable to the expert system implementer wanting to provide authoritative technical Help information.) In desk-bound jobs the tools are ledgers, databases etc. However some tools, which we call "props", are precisely dedicated to assist the expert's problem solving role. We have encountered many such problem solving props in our knowledge acquisition and these are noted as part of the reports in Chapter 3. Consideration of the use of problem solving props has often been very important in understanding the problem solver's knowledge and, for that reason, is reported here.

4.5 · CONCLUSION

The development of techniques has been a supportive exercise for the main body of this research. For this reason, the techniques reported above and the observations accompanying them are not claimed to be comprehensive or in any way particularly ordered. Nevertheless, the reports are all based on very hard experience which, in many cases, has led to the obtaining of results or the development of systems that would have not otherwise been possible. The author very much appreciates that some of his work has been overtaken by elements of the KADS methodology. For this reason, techniques developed for the analysis, structuring and "design" of expert systems are not reported here. However, experience reported in this chapter on the key exploratory stages of knowledge acquisition seems to complement other later work and may prove useful to other workers.

CHAPTER 5 - PROBLEM SOLVING COMPONENTS

5.1   INTRODUCTION

This chapter summarises and discusses the problem solving components that were recognised as features of the problem solving architectures of the occupational problem solving tasks described in Chapter 3.   Appendix B provides a cross-reference chart of the various tasks together with the problem solving components of which we recognise that they are comprised.

The detailed task investigations described in Chapter 3 lead us to believe that occupational problem solving usually involves up to four separate phases. We have termed these the Pre-Search Phase, the Search Phase, the Conclusion Phase and the Feedback Phase.   The Search Phase and Conclusion Phase are considered always to be present but the Pre-Search Phase and Feedback Phase are only present in certain types of problem solving.   Each phase, where it is present, may consist of one or more components.   In many cases such components are alternatives and will be present on an either/or basis.

The components and their functions are considered in relation to the task investigations.   Each problem solving phase will be considered in turn together with the possible components which are present within it.

5.2   PRE-SEARCH PHASE COMPONENTS

Three possible strategic components, Global Data Acquisition, Initial Problem

Analysis and Search Pre-Constraining can precede the actual domain search. Search Pre-Constraining can arise as a result of "Institutional" or "Domain" effects. All three components serve the purpose of increasing the feasibility and efficiency of the subsequent search phase. In particular Initial Problem Analysis and Search Pre-Constraining usually have the function of making a large or intractable problem amenable to solution.

Although no example has been found there is no reason why all three components cannot co-exist. Multinational Tax Planning expertise is somewhat individual to each expert and, perhaps, if there had been an opportunity to investigate a wide sample of experts, we may have found all three Pre-Search Phase Components together.

Initial Problem Analysis was not recognised in our investigations but would seem to be a possibility for the more mature specialist. We have noticed that Global Data Acquisition and Initial Problem Analysis can potentially merge into one component in a mature problem solving task, such as Pensions Management, and see no reason why all three components could not begin to coalesce in Multinational Tax Planning.

5.3 GLOBAL DATA ACQUISITION

Occurrence

Global Data Acquisition appears to be a common problem solving component and is seen in Audit Work Planning, Small Business Guidance etc.

288

Brief Description

The routinised collection (without analysis) of that problem data which will be commonly required throughout the solution to the problem.

Discussion

In many problems experts collect data (perhaps by questioning the person consulting them or by accessing computer data or data on a form) as and when they need it during the process of the consultation. If this were not the case, and all problem information were to be sought in advance, the problem solving would become very inefficient involving the laborious collection of much data which later proved to be unnecessary to solve the problem.

However, in quite a large proportion of problems, experts have recognised that there is a relatively small core set of problem data which will inevitably be required to solve the problem and which, in some cases, will be required in several places in its solution. For instance, in providing Small Business Guidance, the expert will constantly have to formulate his advice in respect of the size of the firm, its turnover, trading sector etc. These factors will be required in analyzing financial performance, commenting on marketing strategy and advising on the structure of the personnel plan.

In such cases where "core common data" exists it is quite usual for the expert to collect it at the beginning of the problem consultation. Often a form is used for this purpose (eg Audit Work Planning) and the data is collected as a fixed procedure rather than as a consequence of the course

of the problem solving. In some cases, as explained in the next section, this collection of core or common data is combined with an initial analysis of the problem. In all cases this element of the problem solving process seems to be one of the signs of a process that is mature and is becoming institutionalised.

5.4 INITIAL PROBLEM ANALYSIS

Occurrence

This component is seen in Small Business Guidance.

Brief Description

A shallow "first pass" through the problem to highlight and prioritise the major issues and, perhaps, to make tentative first conclusions.

Discussion

As just mentioned, the collection of global or core data collection is sometimes combined with an initial analysis of the problem. Different experts for Small Business Guidance adopted procedures which either separated or integrated these components. As an instance of the latter case, the most senior expert would undertake acquisition of "routine" information about the company and, at the same time, gain an impression of the main problems and how the subsequent detailed analysis ought to be

290

prioritised.

Initial problem analysis has a primary function to discover the problem areas (or "symptoms") for the case in question. In the case of Small Business Guidance, the business consultant needs to know and, if possible, prioritise the main areas for attention and guidance before commencing a thorough business analysis. Acquisition of such early "consultation goals" serves to focus and, in most cases, re-order the subsequent search.

In some cases the initial problem analysis also allows a "first pass" problem solution and the subsequent main search then tends to become a rigorous check and elaboration of this tentative conclusion. This was sometimes the case with Pensions Management where a skilled Pensions Officer would be fairly sure of the advice he would have to give just on the basis of the employee's outline description of the problem. However, he would not comment at that stage and would conduct a detailed "semi-confirmatory" session with the employee before making an authoritative statement.

In such cases where there is potential for a "first pass" solution, we can recognise an overlap with the iterative scheme of problem solving discussed later. However, Initial Problem Analysis is characterised by a shallow rather than complete search on the "first pass", often employing different knowledge to that employed on the main search and an emphasis on "problem highlighting" and main search ordering rather than draft problem solution. Also, as in the Pensions Management case discussed above, it is only in a proportion of cases that the Pensions Officer would be able to make a "first pass" solution to the problem. In other cases Initial Problem Analysis would only lead to better ordering of the core part of solving the employee's problem.

Although it is not a task investigated here, we suggest that General Practitioners must often undertake an initial problem analysis (and collection of symptoms) before completing a more thorough medical examination or diagnosis. In some cases the initial problem analysis will lead to referral to a specialist best able to deal with the highlighted problems. In many cases the referred symptoms will include those which the patient has not offered at the beginning of the diagnosis.

Similarly, a small businessman may declare that he has a cashflow problem at the start of a consultation but the initial problem analysis will also highlight that he is also heading for critical distribution problems. The subsequent main search will then be targeted to both the overt cashflow problem and the "discovered" distribution problem.

5.5 SEARCH PRE-CONSTRAINING

Domain search may be artificially "constrained" to a manageable size as a result of "Institutional" or "Domain" reasons. These two forms of pre-search constraint are discussed separately below.

5.6 SEARCH PRE-CONSTRAINING (INSTITUTIONALISED)

Occurrence

This component is seen in Multinational Tax Planning.

Brief Description

Potentially large occupational problems are often tackled on a constrained localised basis to make them manageable. Usually the possible adverse effects of such a "blinkered" approach are counteracted by later feedback to check the global consequences of the constrained decision.

Discussion

Occupational problem solvers regularly and reliably tackle problems which seem unmanageable to the lay outsider. Often the secret of their success lies in the arrangements that their organisations have made to ease the problem solving process. Usually such arrangements seem to have arisen more by "custom and practice" (ie institutionalisation) rather than by design. Nevertheless they are effective. In Chapter 4 we have already discussed the use of "problem solving props" to discipline and support the problem solving process. We here discuss the pre-constraining of problems to make them a manageable size.

Where they have a problem involving many variables, mathematicians often choose to assume that some are held constant in order to achieve an approximate problem solution appropriate to the key variables which interest them. In the same way some business problems, although they are not strictly isolable from the rest of the business, are offered a localised solution which, initially at least, takes little account of any global business consequences of the localised solution.

Multinational Tax Planning provides the only Chapter 3 example of

293

Institutional Pre-Constraining. However such problem constraining by localisation is also illustrated by the task of Corporate Cash Management, a task which we have partially investigated but which is not dealt with as one of the Chapter 3 tasks in this thesis.

Cash management decision making in large companies has to be undertaken quickly, often to get an end-of-day balance and, as a result, management instruct their staff only to consider balances in the local plant or territory. On the positive side this practice allows distributed financial problem solvers to achieve a balance without endless consultation with their counterparties in other parts of the organisation. On the negative side it leads companies to have an unused (and thus wasted) cash balance in one part of the organisation whilst, at the same time, the organisation is borrowing money at heavy short-term rates. (In passing the author suggests that Corporate Cash Management is one of the great unexploited financial expert systems opportunities.)

The institutionalised pre-constrained search in Multinational Tax Planning is somewhat different to that seen in Corporate Cash Management. In Multinational Tax Planning only one problem solver is involved whereas in Corporate Cash Management several problem solvers are involved. This is just one example justifying the further development of what we refer to as "communicative experts systems" which are further discussed in Chapter 6.

In the case of Multinational Tax Planning the pre-constrained search is followed by Global Consequences/Iterative Solution Feedback (discussed below) where the effect of the "local solution" is considered in respect of its world-wide implications. If such implications of the constrained search are markedly unfavourable the problem is again considered on a

294

pre-constrained localised basis but with the further information on the specific factors which made the first solution unacceptable. Thus the possible dangers of the pre-constraining strategy are largely eliminated.

A similar "corrective" feedback loop does not usually occur with Corporate Cash Management largely because time and distribution of the problem between many problem solvers would make feedback almost impossible. However, only on the basis of secondhand evidence, the author has been told that senior managers responsible for cash management in large corporate groups often engage in multiple phone calls until a "good overall" situation is established. Again it must be stated that, with such a poorly optimised process dealing with huge sums of money, the scope for effective expert systems is enormous.

5.7   SEARCH PRE-CONSTRAINING (DOMAIN-DRIVEN)

Occurrence

This very common problem solving component is seen in   Engineering Preventative Maintenance and a number of other common tasks as discussed on the next page.

Brief Description

Sometimes an expert experienced over a wide domain can completely localise his search to a particular sub-domain on the basis of the initial problem description or "at a glance" at the commencement of the search.

Discussion

For the mechanical diagnostic engineer or the medical practitioner their domain of expertise is the "whole machine" or "whole body" respectively. However if the nature of the symptoms make it immediately obvious that just one part of the machine or body should be considered then we suggest that the search is constrained to just that part of the sub-domain. Thus, for the medical practitioner, an initial symptom of breathing difficulties effectively confines the search to the cardio-respiratory system.

In Engineering Preventative Maintenance the initial symptom of a defective product (ie cigarette) with a bruised filling at one end can only be the result of the cigarette "Maker" machine rather than either the machine which makes endless cigarette "rods" or the machine that packs the product in boxes and cases. Thus, at a glance, the problem is totally pre-constrained to the Maker machine. (In the "GP" case quoted above the pre-constraint is not be absolute since breathing difficulties could, in rare circumstances, result from poisoning, psychiatric disturbance etc.)

In Eurobond Warrant Trading the domain is pre-constrained as a result of the particular bond that is to be traded. Traders have expertise, contacts etc across a wide range of business sectors but if, say, a manufacturing bond is required to be traded the domain is immediately constrained to that sub-domain concerned with manufacturing plus the core sub-domains concerned with Reuters interpretation etc. No other sub-domains, such as retail or defence can possibly be of relevance.

In effect experts familiar with this sort of domain are really experts in a number of separate sub-domains. They obviously also have experience at

some whole domain problems such as those concerned with power supplies to the machine, vibration problems etc. (Such problems can be considered to be dealt with by the strategic component "Limited Overarching Conclusions" discussed below.) However their main expertise can be broken down to a small number of sub-domains. We have already quoted the case of medical practitioners who are experienced in a set of sensibly isolable sub-domains (ENT, genito-urinary, psychiatric etc) plus whole domain conditions such as cancer, AIDS etc. Similarly car mechanics can completely separate their "engine noise" expertise from their "radio noise" expertise. Such experts (ie the large machine diagnostician, the GP and the car mechanic) appear to have a very wide span of competence. Nevertheless, in many cases, they need only exercise their competence on a very constrained part of the overall domain.

Pre-constraining of the domain can also be considered from the point of view of a domain consisting of multiple sub-domains (see below) where only one sub-domain must be searched and where there are no reference linkings.

5.8   SEARCH PHASE COMPONENTS

The search phase is the principal phase of any problem solving task. It is during this phase that the domain is searched in order to infer findings in response to the presented problems. In the following paragraphs we consider how the search space may be either considered to represent a simple domain or set of sub-domains. We discuss the nature of search and some of its simpler consequences, such as search space pruning and traversals.

We continue with a major discussion of the nature of search in multiple sub-domains, which may be either a fixed or variable set and which may be reference linked, either minimally or heavily, either to support the problem solving process or to strengthen the final hypotheses resulting from the problem solving.

5.9   SIMPLE DOMAINS AND MULTIPLE SUB-DOMAINS

We recognise two distinct forms of problem domain.   First, there is the simple domain where the knowledge to solve the problem is essentially continuous and homogeneous.   By this we mean that, unless a very fine grain size is taken, it is impossible to split the knowledge into ordered "self-contained" divisions".   Thus in Benchtop Electronic Diagnosis all the knowledge concerns hypothesising and testing in an electronics equipment domain.   In contrast the simple domain approach might not be suitable if the knowledge could be divided into "hardware tests" and "software tests" as may be appropriate for troubleshooting electronic computers.   Both forms of knowledge might be called upon to find a fault but, in most cases, they would be called on serially rather than in parallel.

We first discuss certain features of search in simple domains.   These concern the arrangement of the search space, the site of tests and pruning and transitions through the search space.   Such features also occur in search within the sub-domains of a domain divided into multiple sub-domains.   However, such features are considered here since search within a sub-domain is little different (apart from external references to and from other sub-domains) from search within a simple domain.

5.10 SEARCH

We use the word "search" in a rather universal sense. We see an ordered search space which may be traversed (or searched) in several different ways. Although "search" is an over-used and rather general word we feel comfortable in using it subject to the caveat just given.

The default search is straightforward depth first search and this is the sort of search that, say, a novice diagnostician would undertake on an unfamiliar piece of equipment. Such default search can undergo certain elaborations as experience increases. Such elaborations make laborious depth first search more efficient. We discuss testing at intermediate nodes, local and distant search space pruning and dynamic search space rearrangements and traversals below.

In this section, although there are well-established implementation methods for supporting the described search, we have been particularly careful to work on the basis of observed search behaviour. We have thus attempted to avoid the considerable temptation to attempt to "find" or "map" problem solving behaviour onto some of the very well-established expert systems inferencing paradigms (eg forward or backward chaining).

5.11 TESTS AT INTERMEDIATE NODES

Occurrence

This problem solving component is seen most simply in Commercial Business Underwriting and Foreign Exchange Exposure and in all of the more complex

tasks.

Brief Description

Most domains allow hypotheses to be both generated and tested at intermediate nodes in the search space.

Discussion

A search space can be envisaged where the nature of the domain makes it possible only to undertake tests at the "leaves" even though intermediate nodes in the search space exist. In such a search space intermediate nodes represent sites where hypotheses can be generated but tests are not possible.

In most search spaces intermediate nodes can support both hypothesis generation and hypothesis testing. Tests at intermediate nodes are the route to the search space pruning and traversals discussed below.

5.12   LOCAL SEARCH SPACE PRUNING

Occurrence

This problem solving component is seen most simply in Commercial Business Underwriting and Foreign Exchange Exposure and in most of the more

complex tasks.

Brief Description

Local pruning occurs when a test at an intermediate node is capable of pruning that part of the search space which would have been searched immediately following the test.

Discussion

A test at an intermediate node may be used to check whether or not the search tree need to be searched to a greater depth below that node. Thus, in Commercial Business Underwriting a test for the type of business to be insured will, in some cases prune away the need to establish the requirements for cover for that business.

Thus, if the test for business type shows it to be an Abattoir, there will be no need to establish "Contract Works" cover since abattoirs do not have those risks normally associated with building sites, tunnelling etc.

Such pruning is regarded as local to the intermediate node since, had it not taken place, the search would have moved immediately to the pruned area.

5.13   DISTANT SEARCH SPACE PRUNING

Occurrence

This problem solving component is seen most simply in Commercial Business Underwriting, Foreign Exchange Exposure and Benchtop Electronic Diagnosis and throughout most of the more complex tasks.

Brief Description

Distant pruning occurs when a test at an intermediate node is capable of pruning that part of the search space remote from that which would have been searched immediately following the test.

Discussion

In Foreign Exchange Exposure the expert, in the later stages of providing a means of covering foreign exchange exposure risks, will explore the possibility of asking a potential future supplier to invoice in Sterling or, failing this, US Dollars. No hedging is required for Sterling payments and hedging against US Dollar exchange rate exposures is relatively simple. However, in the earlier stages of problem solving it is necessary to determine the actual currency of the supplier's country. This is required to determine the possible exposure risk. However, as a side effect, if the currency is one of the former Soviet Bloc (or most African) currencies then it can be reliably assumed that the supplier will insist on payment in US Dollars rather than Sterling. Thus the expert, when he actually comes to

devise the hedging strategy, will not explore the Sterling payment option but will move immediately to hedging against future payments in US Dollars.

Thus, in effect, the earlier test for "currency" has, as a side-effect, pruned away the distant "Sterling payment" part of the search space. The effect is said to be "distant pruning" since the "Sterling payment" search would not have immediately followed the "currency" test.

5.14   DYNAMIC SEARCH SPACE REARRANGEMENTS AND TRAVERSALS

Occurrence

This problem solving component is seen in Benchtop Electronic Diagnosis and Small Business Guidance.

Brief Description

Sometimes clues that a problem may exist in a distant part of the search space are discovered as part of routine search. Depending on the strength of such clues, search may transfer temporarily or permanently to the distant part of the search space.

Discussion

In electronic circuit diagnosis, tests for faults in the mains smoothing

circuits are normally undertaken late in the search. Problems which can be assigned to such faults are unusual and it is for this reason that tests are normally only invoked when most more routine problem conditions have been eliminated.

However, it is not unusual for oscilloscope tests on other "more probable" circuit elements to show up problems in mains smoothing. In such case the current tests are suspended and the search moved to the "mains smoothing" part of the search space. Depending on the nature of the triggering condition such search space traversals can be temporary or can result in search space rearrangements.

A temporary traversal is usually undertaken when the triggering condition (eg the indication of something wrong in the distant area) is not a positive indication of a problem in the mains smoothing area of the search space. In such cases it is usual for the problem solver to "interrupt" the normal search routine to "check out" the distant part of the search space. If the problem is isolated to that part of the search space, usually by a test at a distant intermediate node, the search space is rearranged so that the distant area is searched immediately.

If the problem cannot be isolated to the distant part of the search space then search is resumed at its previous place. If the triggering condition is positive (ie a definite indicator of trouble in the distant area) then search space rearrangement takes place without a confirmatory test at the distant intermediate node.

5.15   MULTIPLE SUB-DOMAINS SIMPLE SEARCH

Occurrence

This common problem solving component is seen in Audit Work Planning, Business Performance Assessment, Permanent Health Insurance, Small Business Guidance and Job Shop Scheduling.

Brief Description

The ability to divide a problem solving domain into a number of sub-domains and serially search them (and perhaps combine their results later) is a powerful strategy to cope with the problem solving limitations of the human problem solver. It also leads to a convenient way of breaking down the problem into logical sections for the purposes of asking questions to the "user" of the human expert or expert system.

Discussion

In many problem tasks there is the requirement to consider several different aspects of the problem domain. A search is conducted within each of these sub-domains either to produce a set of primary results (with different results coming from each of the sub-domains) or to produce primary results which lead to lesser or greater global conclusions for the whole domain. The search of the sub-domains may be separate (as in Permanent Health Insurance where "Health" and "Occupational Risk" are

totally unconnected searches) or may be linked in one or more of the various ways discussed below.

Dealing with a problem domain by way of a number of essentially separated sub-domain searches is one of the most powerful devices in occupational problem solving. Lay problem solvers are often impressed at the size of the domain which an expert can command. However, experience within the organisation and within the individual often means that an apparently large unmanageable domain can be split up into a number of smaller sub-domains which can be processed serially.

Each of these serial searches is such that it can be coped with by the problem solver and, once completed, it is often only necessary for the "result" to be recorded before proceeding to the next sub-domain.

5.16   MULTIPLE SUB-DOMAINS - FIXED SET

Occurrence

This problem solving component is seen in all tasks involving Multiple Sub-Domains except Eurobond Trading.

Brief Description

A fixed set of problem solving sub-domains is said to exist where the domain can be divided into a number of sub-domains each of which must be

searched (with certain well-defined exceptions) whenever a problem is presented.

Discussion

In most cases it is possible to determine the set of sub-domains which are considered in order to give an overall problem solution. This "Fixed Set" may actually vary from problem type to problem type within the expert's repertoire but, for any one problem type, the expert will search a fixed set of sub-domains. Thus for Small Business Guidance the fixed set is Financial, Purchasing, Production, Distribution, Marketing, Personnel and Management. No business can be fully advised unless all these sub-domains are searched (unless, for instance, it is a service organisation where Production and Distribution do not apply) and no business requires any further sub-domains to be searched. In such case the set of sub-domains is fixed and set at six, rather than seven, at the outset of the problem.

5.17   MULTIPLE SUB-DOMAINS - VARIABLE SET

Occurrence

This problem solving component is seen in Eurobond Trading.

Brief Description

A variable set of problem solving sub-domains is said to exist where the

domain can be divided into a number of sub-domains each or any of which may, or may not, be searched depending on the nature of the problem or, more particularly, the progress of the problem solving. Such "variable set" problems are particularly difficult to solve.

Discussion

The single example Eurobond Trading illustrates a component of problem solving which makes problems very difficult to "solve" on any sort of reliable basis. In essence the problem solver must consider the strategic question "What sort of sub-domains should I be exploring to solve this problem" concurrently with the actual problem solving activity itself.

Once a set of solutions has been established from the selected sub-domains the expert usually combines their findings to a single problem solution. In the case considered, Eurobond Trading it was clear that the expert does not approach problem solving in two stages by first establishing the set of sub-domains and then exploring them. In contrast sub-domains are dynamically entered and dropped as the problem progresses.

The Eurobond Warrant Trader knows many possible sub-domains that may be explored to set a price. Such sub-domains are truly separate and different knowledge is used to progress problem solving within them. Some of the sub-domains are more-or-less obligatory, such as those concerned with consulting Reuters screens or seeking quotations from market counterparties. Others are less common, such as investigating the accounts of firm's in the same market as the Bond Issuer or using one of many Decision Support Tools. The lack of a fixed set of sub-domains to search

308

makes problem solving not only difficult but, since "critical" domains may be easily missed, very unreliable.

It is necessary to make clear the difference between the fixed/variable sub-domain dimension and the open/closed problem dimension. "Open problems" (ie problems without a finite set of possible solutions) are, of course, also very difficult to solve. However a variable sub-domain problem may have a closed solution set. In the case of the Eurobond Warrant trader the solution set may be just the binary "Buy" or "Sell" set. Nevertheless the "openness" of the sub-domains to be considered makes the problem just as hard as that resulting from the "openness" of the solution set in other problems.

Some problems, such as those in the military and political domain, would seem to be "open" (ie with infinite possible solutions) and also to lack a fixed sub-domain set for problem solving. Solutions to such problems are often considerably sub-optimal even when the problems are tackled by the most knowledgeable and intelligent experts.

5.18   SUB-DOMAIN REFERENCE LINKING

We have already noted that the ability to treat a large problem as a the serial search of a number of sub-problems is a powerful experiential route to making occupational problems tractable. A certain amount of "interrupts" to any of the sub-domain searches to reference other sub-domains is tolerable. Nevertheless problem solving becomes difficult and error-prone where this linking effectively means that the serial sub-domain searches coalesce to a parallel search of all the sub-domains constituting the overall

domain.

In the following paragraphs we consider both "minimal" and "heavy" linkage between sub-domain searches and linkages for the purposes of "problem solving" and/or "hypothesis strengthening".

5.19 MULTIPLE SUB-DOMAINS – MINIMAL REFERENCE LINKING SEARCH

Occurrence

This problem solving component is seen in Small Business Guidance and Telecom Repair Service.

Brief Description

Sub-domains may be serially searched reasonably efficiently in those cases where references to other sub-domains are minimal. We suggest that two references from any one sub-domain to another represents the upper limit of "minimal reference linking".

Discussion

The completion of the search of a particular sub-domain may require information from three sources. First, there may be global data acquired at the beginning of the problem which is relevant to the sub-domain under

consideration. Second there may be data uniquely proper only to the searched sub-domain which must be acquired directly at the time of problem solving. In Small Business Guidance, the small businessman may be asked about capacity of distribution vehicles as and when needed during consideration of "Distribution" sub-domain issues. This information is sufficiently rarely required not to be included in the initial collection of core global data. As a result it will be acquired by "reference to the user" during the search of the Distribution sub-domain.

The third category of information, which is the subject of our interest in this section, is derived data required from other sub-domains. Such derived data may be readily available as a partial result of a prior search of the other sub-domain or, if the sub-domain has not yet been searched, from a special partial search to support the current focus-of-attention sub-domain. Thus, rather than "reference to the user", there is "reference to another sub-domain". Naturally, this could actually involve an indirect "reference to the user".

For instance in Telecom Repair Service the Diagnostic Testing sub-domain can often provide a complete diagnosis of a remote line fault. However, in certain cases the Diagnostic Questioning sub-domain must be used so that the customer can provide information needed for the Diagnostic Testing. For instance, if there is a load hum on the line this can be caused by mains interference either in the underground cables or at the customers premises. Diagnostic Testing would proceed to isolate the fault up to the point where it was necessary to question the customer about the internal cables at his/her home. Long cable runs in the proximity of mains cabling would be a sure indication of interference inside the customers premises and would enable the expert to return and continue Diagnostic Testing to an early

positive conclusion.

We suggest that minimal reference linking may be considered to exist when not more than two linkages, either for problem solving or for hypothesis strengthening, are made in any one direction between two sub-domains. In other words we feel that "minimal" means that no one sub-domain references another more than twice in the solution of a particular problem. In most of what we have seen, in Telecom Repair Service and Small Business Guidance and in other unreported tasks, there usually exist only single linkages between sub-domains with two linkages being quite exceptional. Therefore our figure of "two or less" seems a reasonable guideline for the purpose of definition.

5.20   MINIMAL SUB-DOMAIN REFERENCE LINKING (PROBLEM SOLVING)

Occurrence

This problem solving component is seen in Small Business Guidance and Telecom Repair Service.

Brief Description

Inter-sub-domain references to advance the progress of the problem solving rather than to strengthen solution hypotheses.

Discussion

Most of the sub-domain linkages we have encountered are for the purposes of providing information to actually allow each sub-domain search to continue to its conclusion. We refer to such linkages as "Problem Solving" linkages to differentiate them from "Hypothesis Strengthening" linkages discussed in the next section.

5.21 MINIMAL SUB-DOMAIN REFERENCE LINKING (HYPOTHESIS STRENGTHENING)

Occurrence

This component is seen in Telecom Repair Service.

Brief Description

Inter-sub-domain references to strengthen or "confirm" solution hypotheses rather than directly advance the progress of the problem solving.

Discussion

Sometimes a single physical domain can be amenable to analysis by two distinct sets of knowledge. We have taken to referring to this situation by the term "overlapping sub-domains" because the "domain competence" of the two sets of sub-domain knowledge is not identical but has some areas of

313

common applicability.

To clarify this statement we refer to Telecom Repair Service where both Diagnostic Questioning of the customer and Diagnostic Testing on the line can, in some cases, equally and independently give rise to a diagnosis of "Earth Fault" on one the customer's line. This is the "area of overlap". However Diagnostic Questioning is the sole way to isolate certain faults in PABXs and Diagnostic Testing is the only way to diagnose the location of many underground faults. These are areas of "non-overlap".

In Telecom Repair Service there is no absolute proof that a correct problem solution has been established. The diagnosis will only be proved correct when a faultsman actually goes out and checks a distant PABX or underground cable. (Similarly in Life Underwriting the correctness of the underwriter's "diagnosis" will only be proven when the insured person dies at, or before, the actuarially anticipated age.)

Since the consequential cost of errors in such "diagnostic hypotheses" is high it is often the practice of experts to use the "alternative" knowledge to "confirm" (or, more properly, "strengthen") the diagnostic hypothesis. Thus, in Telecom Repair Service, it is common for an Earth Fault suggested by Diagnostic Testing to be followed by a call to the customer in which careful Diagnostic Questioning will attempt to provide a "parallel confirmation" of the hypothesised Earth Fault.

Such hypothesis strengthening reference linkings between sub-domains (even if they are overlapping) are very similar in form to problem solving reference linkings and we again suggest that these can be regarded as minimal if no one sub-domain references another more than twice in the

314

solution of a particular problem.

5.22 MULTIPLE SUB-DOMAINS - HEAVY REFERENCE LINKING SEARCH

Occurrence

This problem solving component is seen in Job Shop Scheduling and Eurobond Trading.

Brief Description

Heavy referencing between separate sub-domains, either to support problem solving or hypothesis strengthening, largely offsets any advantages that the problem solver may gain by conducting serial sub-domain search.

Discussion

Both Job Shop Scheduling and Eurobond Trading demonstrate dense reference linking between sets of knowledge which, although they have a large domain overlap, can be largely regarded as separate. In the case of Job Shop Scheduling the references are problem solving references largely concerned with constraints and constraint relaxations. In the case of Eurobond Trading the references are primarily hypothesis strengthening references concerned with confirming the correctness of an offered or bid

315

price or decision to buy or sell.

In both cases the denseness of the references seems to largely counteract the advantages accruing from serial sub-domain search. As a result the problems become intractable for average problem solvers and; accordingly, the optimality of solutions is relatively low.

5.23 CONCLUSION PHASE COMPONENTS

Where there is no division of the domain into separate sub-domains we consider all conclusions to a consultation to be "primary". Where several sub-domains are searched there is the possibility that the primary conclusions from each of the sub-domains may, to a greater or lesser extent, be combined to, supplemented with or replaced by secondary conclusions. There may also be some limited aspects of the domain which cannot be considered on a compartmentalised sub-domain basis and which lead to conclusions over and above the predominant primary and secondary conclusions. In the following paragraphs we consider the various forms of conclusion components arising in a multiple sub-domain problem.

5.24 SECONDARY GRAND CONCLUSIONS

Occurrence

This component is seen in Business Performance Assessment.

Brief Description

Although each sub-domain may have its own useful primary conclusions it is often the secondary conclusions which are most valuable. Such secondary conclusions are based on the "summation" of the primary conclusions and are not obtainable directly from the problem symptoms.

Discussion

In some types of problem the separate analyses of a number of sub-domains can all contribute to important "feelings" or "summations" about underlying problems affecting the whole domain. For instance, in Business Performance Assessment there are very separate analyses of the performance of a number of different business activities. However, all these business activities are governed by the same business policy and their analyzed performances all contribute to the understanding and assessment of the policy. The policy cannot be accessed or measured directly so it must be evaluated on the summation of the results of a number of parallel independent measures on the activities which result from its implementation.

Where this form of problem solving component exists, it is usually the secondary "grand" conclusion which is the most valuable outcome of the consultation. Nevertheless, the individual primary conclusions are essential in their own right since they reflect detailed issues proper only to their sub-domain. They may also show local aberrations which are contrary to the overall picture of the problem. Thus, in Business Performance Assessment, a hospital may be evaluated (as a Secondary Grand Conclusion) as having excessive spending on facilities at the expense of proper staffing.

317

Nevertheless, as a particular primary conclusion, it may be shown, in its geriatric unit, to be excessively staffed at the junior level as a result of poor specialised facilities to cope with needs of old people. This primary conclusion is in direct conflict with the conclusion about the hospital as a whole.

5.25   MASSIVE COMBINATORIAL CONCLUSION MATRIX

Occurrence

This component is seen in Permanent Health Insurance.

Brief Description

The shallow analyses of a number of sub-domains can result in a potentially massive number of possible primary conclusion combinations. Sometimes the resultant secondary conclusions are generated laboriously, and somewhat unreliably, for each of the combinations.

Discussion

We have already discussed global data collection and initial problem analysis as separate topics above. However there is a category of "primary search" which is little more than initial data collection. It is then followed by secondary "massive combinatorial" conclusions which we discuss below.

Permanent Health Insurance is the particular example and we are aware of other tasks which, although we have not investigated them in detail as with Permanent Health Insurance, we are convinced would provide equally good examples. In such tasks a number of sub-domains (often more than five or six) are analyzed at a fairly trivial level (ie just one stage removed from the collection of raw data) to produce a string of simple primary conclusions. To quote from Permanent Health Insurance, the country of residence of the applicant for insurance will be analyzed into just three categories - Normal, Risk and Extreme Risk. (As examples both Korea and Mexico would be in the "Risk" category, the former due to risk from terrorism or warfare and the second because normally occurring medical conditions would be unlikely to receive adequate treatment. Iran and Pakistan would be in the "High Risk" category for similar, but more extreme, reasons.)

The number of possible secondary conclusions from combinations of the string of primary findings is extremely large and, in our investigation of Permanent Health Insurance, we found that each set of combinations was dealt with from first principles. This process was extremely time-consuming and, since much guesswork and judgement was involved, there were many anomalies in the established conclusions. Previous conclusions were recorded and common sets of combinations were well known. For instances most of the cases where all factors were "normal" (ie healthy, UK, non-dangerous occupation etc) were simply and reliably judged on age and level of income assured.

As a result, provided that the primary conclusions can be established and the combinatorial rules established, the "combinatorial conclusion" component of occupational problem solving is one of the most suitable for productive

319

expert systems support.

5.26 LIMITED OVERARCHING CONCLUSIONS

Occurrence

This component is seen in Audit Work Planning.

Brief Description

Even where a domain can be analyzed as a set of separate sub-domains there will usually be a limited number of conclusions (or solutions) which can only be found by considering factors from the domain as a whole. Sometimes such conclusions offer the only problem solution.

Discussion

In crude terms "Limited Overarching Conclusions" are those conclusions which are left over when everything else has been tidily dealt with on a serial sub-domain basis. Such overarching conclusions are not necessarily any more important than the primary conclusions found in the sub-domain searches. In some cases they may be based on a mixture of primary problem data and primary conclusions but they are not really secondary conclusions. They are just the untidy issues remaining when the rest of the domain has been analyzed in an orderly compartmentalised basis.

320

In our discussion of Audit Work Planning we have already provided a good example of such overarching conclusions. We have also encountered overarching conclusions in the analysis of certain electronic equipment diagnostic tasks where a fault has not been found despite the separate and meticulous checking of all the circuit sub-assemblies. At this stage the technician has concluded that, although each sub-assembly is "within specification", there may be an interaction between sub-assemblies which are separately at "opposite ends" of their specification tolerance. (For instance, in some work on Pulse Code Modulation equipment, timing tolerances of coder and decoder boards were both within the specified limits but the technician had noticed that the coder board was near the lower limit and the decoder board at the higher limit. At this "worst case" situation they would not work satisfactorily with each other.)

In such cases the overarching conclusion becomes the final and vital conclusion to the task whereas in Audit Work Planning such overarching conclusions are of similar (or lesser) status to all other conclusions.

5.27   POST-SEARCH CONSTRAINED CONCLUSION (INSTITUTIONALISED)

Occurrence

This component is seen in Life Underwriting.

Brief Description

Allowing the problem solver to provide a "less accurate" solution than is

actually possible (or offering the option to decline the opportunity of solving the problem) is a common way of increasing problem solving productivity.

Discussion

We have already seen how organisations allow problem solvers to cope with otherwise unmanageable problems by artificially constraining the problem to just part of the domain. In the same way it is possible to ease the problem solvers task by increasing the grain size of the conclusions allowed to him or her. Thus, although a problem solver might struggle and eventually succeed in providing a very "accurate" problem analysis, a quicker and less accurate analysis will often suffice for his organisation.

Permanent Health Insurance and Life Underwriting provide an interesting contrast. In both cases similar details are available to the underwriter (Age, Health, Occupation, Sports etc) and in the former, as we have already discussed, an "accurate" predictive interpretation is attempted. In the latter the requirement is almost reduced to an "Accept" or "Refuse" decision. ("Accept at increased" rates is also used but with only two or three possible increased rates.) In such cases the burden on the decision maker is both reduced by the simplicity of the categorisation and also management's instruction to "fail safe". This the underwriter for Permanent Health Insurance" is obliged to provide a quotation in nearly all cases whereas the Life underwriter is instructed to decline doubtful cases. Speed of problem solving is more important than accuracy (or maximum capture of business) in tasks characterised by institutionally constrained conclusions.

5.28   LOCAL OPPORTUNISTIC CONCLUSIONS

Occurrence

This problem solving component is demonstrated in Engineering Preventative Maintenance.

Brief Description

Following solution to the primary problem the information discovered to enable that solution, or the physical situation presented by the procedures undertaken, may allow further conclusions to be made on an "inexpensive" opportunistic basis.

Discussion

We have only one instance of this component although we feel it likely that it probably exists in similar mechanical diagnostic tasks to Engineering Preventative Maintenance.   We have had only a single opportunity to investigate a "mechanical" task and this is probably why only a single instance of this component has been found.

Although we have informally classified Local Opportunistic Conclusions together with other secondary conclusion phase components following multiple sub-domain search, it could also easily exist as an adjunct to simple single domain search as seen in Commercial Business Underwriting and

Foreign Exchange Exposure.

In Engineering Preventative Maintenance the diagnostician is obliged to disassemble equipment in order to rectify the primary component failure. Similarly, the examination of the mechanisms related to the failed component provides him with immediate visual clues to the abnormal wear or maladjustment of components in the vicinity of the failure. Such opportunities of access and cost-free diagnostic clues prompt the diagnostician to make conclusions as to the need for additional further work. This work is usually then immediately undertaken as part of the overall maintenance intervention.

Analogous examples of Local Opportunistic Conclusions from other domains are easy to demonstrate. For instance the primary activity of an auditor is to explore a company's financial and other records in order to verify their accuracy. However, during this process many other facts are discovered and observations made which can be offered as "value added" additions to the firm undergoing the audit. Such additions essentially comprise an aggregation of Local Opportunistic Conclusion components.

5.29   FEEDBACK PHASE COMPONENTS

In many problem solving situations the "final results" are evaluated, either for their consequences on the overall domain or, as a separate assessment activity by the problem solver. If such "final results" are found to be in some way unsatisfactory, the problem solving cycle is repeated in an attempt to find a better solution.

5.30   LOCAL RESULT GLOBAL FEEDBACK

Occurrence

This problem solving component is seen in Multinational Tax Planning.

Brief Description

"Global evaluation" is an effective and relatively easy way of eliminating any seriously non-optimal consequences arising from institutionally constrained (or localised) problem solving.

Discussion

We have already seen in Multinational Tax Planning that search is restricted to an artificially constrained sub-domain as a means of making the problem tractable to the expert.   The concomitant requirement is to submit any conclusions based on this restricted search to an evaluation based on the effect of the proposed solution on the overall domain.   This is the problem solving component which we refer to as Local Result Global Feedback.   If the effect of the local result has some unexpected negative effect in the global context then the problem solving cycle is re-started in the light of this new information.

Local Result Global Feedback is a very common component of corporate problem solving especially where the problem solving is distributed between a group of problem solvers who may be unaware of the global consequences of their local decisions.   Unfortunately the time delays and lack of rigour

325

applied in such situations often mean that the "negative feedback" often arrives too late or unreliably for it to prevent local decisions having their adverse overall affects.

5.31 FEEDBACK FROM CONCLUSIONS AS INPUT FOR ITERATIVE SOLUTION

Occurrence

This problem solving component is seen in Audit Work Planning.

Brief Description

"Non-scientific" problem solvers can usually recognise a good (or bad) solution more reliably than they can generate one. This component utilises this ability as part of an iterative problem solving cycle.

Discussion

It is often easier to recognise a "good" solution than to generate one. In particular we have noticed that problem solvers from domains which do not have a "scientific method", primarily financial domains, tend to use an iterative method of problem solution to attain a result which satisfies them. "Scientific" problem solvers tend to work through a problem once, applying scientific or quasi-scientific principles, and arrive at a final solution at their first attempt. Any "experimentation" is done within the problem solving by use of embedded "hypothesise and test" methods.

326

In contrast problem solvers from "non-scientific" domains often work completely through a problem and, once an end-finding is established, apply an overall evaluation as to its "correctness". Audit Work Planning showed this problem solving component to a marked degree. Other investigated tasks (eg Small Business Guidance) also showed it to a lesser extent although it has not been specially highlighted in the analyses of these tasks.

Iterative Feedback also plays a particular role where the problem solver has made an initial problem analysis and has perhaps inappropriately focused the problem. On the second iteration, based on the awareness of the first outcome, the problem solver is able to make a far more optimum focusing of the problem.

We suggest that problem solvers may have at least two sorts of knowledge. Primarily they have knowledge about "how" to solve a problem. Secondly they have knowledge about "what" is a good solution. This latter knowledge may be in the form of a repertoire of known good solutions which they can compare with the current tentative solution. A similar form of knowledge has also been recognised in experts for Job Shop Scheduling. Many schedulers are reputed to be able to recognise a "good" schedule even although they have great difficulty, or find it impossible, to generate one.

5.32   TOWARDS A GENERAL MODEL OF PROBLEM COPING

The series of diagrammatic representations of problem solving components built up in Chapter 3 combined with Appendix C would seem to have some

327

additional value in that they might form the beginnings of a rather particular general model of occupational problem solving. Such model could be regarded as orthogonal to other models, such as the KADS interpretation models. Such other models are primarily concerned with the problem solving knowledge associated with the task. In contrast a model based on the strategic level components reported in this thesis is more-or-less wholly concerned with the strategies used by experts to cope with limitations in their own cognitive processes. Non-experts, unaware of such strategies even thought they may have other knowledge (represented in existing models) relevant to the tasks, might well be defeated by them or would undertake them in a very protracted and inefficient fashion.

Following further investigations of occupational problem solving tasks (or the re-visiting of some tasks investigated in the absence of the model), leading to greater confidence in a set of components based on those discussed in this chapter, the model derived from them could be used both to assist and guide both knowledge acquisition and understanding relevant to further occupational tasks.

5.33   PROBLEM SOLVING STRATEGIC LEVEL COMPONENTS: CONCLUSION

This research has led us to recognise and propose a particular set of problem solving strategic level components.   Such components are as described in the foregoing sections of this chapter and summarised in Appendix C. In some cases the proposed components have been recognised on multiple occasions and, in other cases (eg Institutionalised Pre-Constraining) only a single instance is reported.

In discussing the components we have shown how problem solvers are sometimes faced with problems which are so large or complex that they would be unable to cope without problem solving being conducted using the strategic components as described. In particular we note that a primary purpose of the components is to provide the problem solver with a series of small searches, with which they can cope. Where such components cannot be used, and the search or its results cannot be constrained, or where it is impossible to adequately divide up the search, then the unaided problem solver would seem to have severe difficulties in coping with the problems with which he or she is presented.

As we discuss in Chapter 7, it may be that some very worthwhile knowledge-based systems opportunities exist in just those areas where humans are not able to "divide and conquer" the problem. The recognition of the appropriateness of such opportunities could represent an important applied outcome of this research and could not have arisen without an understanding of the strategic components and their limitations reported here.

In an attempt to generalise our findings we must recognise two important limitations to our proposals in respect of problem solving strategic level components.

First, although we have investigated a reasonable selection of occupational tasks, it would require many more task investigations to demonstrate that the components we propose occur generally in occupational problem solving.

Second, it is rather unlikely that the strategic level problem solving components described in the above previous sections form a correct or

exhaustive list. Not only is it very possible that some, or most, of the proposed components require partial re-definition but it is inevitable that there must be further, as yet unproposed components existing in occupational tasks not investigated during the course of this work.

But, setting aside for the moment the issues of the correctness or completeness of the proposed components, what is their academic or commercial value? In the above paragraphs we have above suggested that, following further task investigations, a refined taxonomy of task components along the lines of those proposed here could be used as the basis of a model to support knowledge acquisition and understanding of other occupational problem solving tasks. This issue is explored further in Section 7.3 as part of a commentary on the possible relevance of this research to KADS.

Such an applied usage would be a useful secondary outcome of this research. The problem solving strategic level components summarised in this chapter were not sought as an initial objective to this research but were recognised, and then formalised, as a direct result of our objective to understand how occupational problem solvers coped with complex problems. However, the fact that such objective has led us to begin to formalise a new model of problem solving (or, perhaps more correctly, a model of problem coping) could, we believe, provide a useful tool for both the academic and commercial investigator of occupational problem solving.

CHAPTER 6 – FACTORS FOR EXPERT SYSTEMS FEASIBILITY SUCCESS OR FAILURE

6.1 INTRODUCTION

This chapter offers the last of the three perspectives based on the findings of Chapter 3 and provides a more complete examination of the factors that affect whether a potential expert systems project will move forward to an implementation that will be of business benefit to the sponsor.

The author has been reasonably successful in his commercial expert systems work. Nevertheless, he has investigated many expert systems opportunities which have never proceeded to implementation and, even where a system has been developed, this has not always resulted in its long term usage for the benefit of the sponsor. Using the findings of Chapter 3 this chapter discusses those factors which predispose to the feasibility, success or failure of a potential expert systems project.

The chapter commences with three different dimensions of feasibility. First, is it technically feasible actually to build a proposed expert system? In the main this section concentrates on whether or not the supposed expert knowledge actually exists and can be acquired. Second, is it possible to construct or provide an expert system in the organisational climate of the sponsor? Third, even if a system is feasible and appropriate for the organisation, will it be cost effective to implement and, probably more important, to maintain? This last issue continues to be the key stumbling block for most projects and will be considered in some detail.

Four further specific issues are then considered. They are examined as

special topics because of their importance to overall systems feasibility. The issues concern knowledge acquisition, user interfaces, "expert-user appropriateness, and hardware.

Finally, there appear to be two potentially very fruitful areas of commercial expert systems opportunity which require to be separately identified and discussed. We refer to these as "communicative expert systems" and "corporate knowledge systems" and they form the final part of this chapter before a short concluding summary.

There have been many very able and dedicated workers in the expert systems field. Nevertheless, the sum total of active expert systems in commerce and industry is low. In many cases projects have failed not on the basis of the skill of the implementor nor from limitations in the available functionality of the technology. Other issues are usually responsible for expert systems success or failure and it is the task of this chapter to order and describe those encountered in this research so that they may be duly considered by other workers in advance of expert systems projects.

6.2   FUNDAMENTAL ISSUES OF SYSTEMS FEASIBILITY

This section considers fundamental technical issues which can prohibit the development of an expert system. It will be noticed that "availability of knowledge" is the key issue underlying all except the last of the topics in this section. The author in reviewing this section concludes that, although it has sometimes been difficult or laborious, he has never lacked the technology to implement a system provided that stable, reliable knowledge was available at the outset.

Perhaps this is a salutary lesson to those who would point to the need for smarter and more complex tools as the route to successful commercial expert systems. Where the technology is inadequate it is for much more fundamental reasons, such as the inability to understand continuous speech (as mentioned with reference to Small Business Guidance in the discussion at the end of this section on the need to design "substitute" functionality) rather than the more immediate expert systems reasons.

Large Problem Space and Few Examples or Principles

Permanent Health Insurance provides the classic case of this problem. The potential number of combinations of age, health status, occupation, risk, cover etc is enormous and there was no chance of investigating anything more than a fraction of these cases. It was very likely that, had this been done, the expert would have either not have experienced the case in the past or would not have been able to recall it.

(In anticipating some of the content of Section 6.9 we might suggest that Permanent Health Insurance remains in the province of Knowledge Creator Experts and that such experts, who are primarily solving novel problems, are undertaking the sort of "intelligent" activity evaluated in intelligence tests and assessments. However, expert systems currently are directed at problem solving tasks where the search space has been relatively well explored or the problem solving is based on principles that can be established on knowledge acquisition from the expert.)

In the case of Permanent Health Insurance we attempted to establish the principles upon which the expert worked in order to combine the various

risk factors. (Such principles would, of course, provide the basis of a rule set to deal with risk factor combinations.) However, when a sample set of hypothetical cases was investigated it was impossible to get anything like a reasoned evaluation of risk. Instead the underwriter would claim that the stated case was "something like" another case or would state that further investigations were necessary. It could be argued that this was a justification for a case-based reasoning system. However, the author is far from convinced that this would be useful. In fact, as stated in Chapter 3, he suggests that the underwriter was acting as a "wise and unquestioned" authority who made decisions on a reasonable but fairly arbitrary basis.

In such cases, where the problem space is large and there are few instances and no opportunity to acquire the principles upon which problem solving is conducted, it seems unlikely that an expert system can be built. Potentially, if the sponsor and expert are frank and cooperative, such situations present an opportunity to explicitly establish the principles underlying the task and design a system to incorporate these. However, in so doing, the expert system would have moved away from its normal role of simulating existing expertise and would become equivalent to a conventional system based on the analysis of a business requirement etc. This is a valid use of expert systems technology and is further discussed in Chapter 7.

Lack of Reliable Expertise

In the last case the key problem was "lack of expertise". Even though his knowledge is believed to be very incomplete, the "expert" in Permanent Health Insurance and similar tasks is generally assumed to be competent and his decisions go largely unchallenged. However, a separate problem arises

334

when the expert has expertise but the use of that expertise does not reliably end in a good outcome.

For instance in the case of Eurobond Warrant Trading, it is agreed by the experts themselves that their decisions lead to a favourable financial outcome in only just over 50% of trades made. Thus nearly half of the expert decisions lead to a financial loss.

Of course, it is accepted that an "expert" does not require to be correct on 100% of occasions. However, when the "hit rate" is little more than evens it presents the expert systems builder with an impossible task to build an effective expert system. In the case of Eurobond Warrant Trading there were additional factors, such as the instability of the expertise and the difficulty of encoding verbal information, that made an implementation infeasible.

Nevertheless, the key factor was the reliability of the expertise. In this respect "reliability" refers to the reliability of the outcome rather than the reliability of the expert. Thus, for any given situation, the expert would reliably give the same advice (subject to the "stability" issue discussed immediately following). However, such advice only proved to be reliably correct (ie led to profit) in just over half the cases. Therefore the expert was consistent and "knowledgeable" but his knowledge could not be relied upon to provide a useful outcome. In contrast the expert for Permanent Health Insurance expert lacked knowledge (ie was not knowledgeable) and would have been inconsistent in his risk evaluations.

The financial trading/dealing arena has attracted many knowledge engineers and sponsors who thought that immense gains could be made from expert

systems based on the knowledge of expert traders and dealers. Much money has been wastefully invested and many careers damned by these endeavours. The reputation of expert systems, especially in the financial community, has been damaged by these activities. In retrospect it is difficult to see how any worthwhile expert system could have been built on the basis of such unreliable expertise.

Lack of Stable Expertise

Eurobond Warrant Trading provides us with a further pitfall to successful expert systems implementation. Even if the expertise exists (as may not have been the case with Permanent Health Insurance), and it is reliably correct and effective (as was not the case with Eurobond Warrant Trading), there is little hope of implementing a worthwhile expert system if the underlying knowledge is changing or developing from day to day.

Financial traders, politicians and, perhaps, economists seem to have what they call a "view" of how their domain works. Certainly in the case of Eurobond Warrant Trading the trader admitted that this "view" would change and develop on the basis of his experience of the market and its development (including the changes and developments in other traders). The "view" is the basis on which the trader operates and, effectively, constitutes his current set of heuristics as to how to make money from the markets.

Apart from the difficulty of actually capturing the trader's expertise, which was at best unreliable, the fact that such expertise was continually changing would have made an expert systems implementation totally impossible.

Lack of stable expertise is also a factor raised under "Economic Viability" below in respect of Benchtop Electronic Diagnosis. However, in such latter case there was good reason for the knowledge to change as a result of actual circuit modifications. In the case of Eurobond Warrant Trading the "view" would change with very much less justification and perhaps largely on the basis of the trader's experimentation to determine whether or not the trading success rate could be increased by changing his "view".

Requirement to Design "Substitute" Functionality

It is not always possible to reproduce the style of working undertaken by the expert and, in those cases, successful implementation or otherwise will depend on the ability to design "substitute" functionality for that which cannot be reproduced.

A good example is provided by Small Business Guidance where one of the key skills of the expert business consultant was to listen to the outpourings of the businessman when asked about his business. The interpretation of this torrent of material enabled the expert consultant to collect considerable material upon which to begin the more structured part of the session with the businessman. (Such initial phase constitutes the Global Data Acquisition and Initial Problem Analysis components of the expert task.)

This problem was actually overcome in the implementation of the Small Business Guidance expert system. As detailed in Chapter 3 an initial supportive dialogue was conducted with the businessman user in order to achieve the same objectives as those of the human consultant listening to the businessman. Nevertheless, to be clear, the designed "front end" of the

system with its engaging route maps and feedback provided a user interaction which was very different from that offered by the human consultant. Therefore the success of the system depended very much on the ability to design substitute functionality to that which could not be offered by the system.

Although it was possible to substitute for a "listening consultant" in Small Business Guidance system the author could not achieve a similar substitution in Eurobond Warrant Trading. Traders are regularly in conversation and any one individual trader is constantly changing his position on the basis of news and gossip acquired from other traders. There is no way that any currently existing speech understanding system could deal with such verbal input. Similarly, however ingenious and rewarding its design, it would be impossible to force the trader to undertake an artificial dialogue with the expert system so that he could encode the information received from other traders. Therefore, in addition to its several other problems, Eurobond Warrant Trading was not implemented because a vital part of its functionality could neither be simulated nor substituted.

It is interesting that Job Shop Scheduling suffered a similar problem and it was not possible to design substitute functionality. In this case the "verbal work" would have been undertaken by the system's user which would have made the overall system uneconomic to implement.

6.3   ISSUES OF ORGANISATIONAL FEASIBILITY

Fundamental technical issues prevent a system being implemented however strong the organisational goodwill or funding. In contrast, organisational

338

(and economic) issues can prevent a project taking place and an expert systems implementation being made and deployed. Once recognised, organisational issues can sometimes be countered as with Life Underwriting. In other cases, such as Engineering Preventative Maintenance, this is not possible and an otherwise valuable project fails to mature. Key issues of organisational feasibility are discussed in the following paragraphs.

Experts, Sponsors and Users Addressing Different Problems or Priorities

Expert/sponsor mismatch in priorities was detected by the author in Life Underwriting where senior underwriters (ie managers) were keen to refuse or load the premiums of the minority of applicants declaring potentially serious risks (eg dangerous "non-hobby" sports, alcoholism, history of tumours etc) but, in the main, were not interested in the minor risks and variations presented by the bulk of the applicants. In the latter case they were keen to process as many applicants as possible at normal rates and get them accepted so that premiums could start flowing into the company.

In contrast the training of underwriters used as experts for knowledge acquisition led them to believe that they should make the most accurate prediction for all possible applicants. In fact, in a rather childish detective fashion, they enjoyed discovering linked facts about an applicant that could lead to an early death (eg the potentially suicidal "thin teacher type" mentioned Chapter 3). Such matters were of no interest to the more senior underwriters who realised that an allowance for a proportion of such cases was already covered in the standard rates offered to "normal" applicants.

Once this disparity between expert and sponsor had been detected the

investigation and implementation proceeded to a very successful conclusion. The implemented system simulated the expert's knowledge but the goals of the system were skewed towards management's desire for throughput rather than the expert's desire for accuracy.

In Business Performance Assessment there was a mismatch between the eventual users' priorities and those of the initial sponsor and expert. Fortunately this mismatch, only realised after the implementation, was beneficial to all concerned. As reported in Chapter 3 the goal of the system was to produce reports for senior management so that they could undertake regular management appraisals of operational units. However, the system once implemented and deployed became used by those operational units to improve their own performance before problems needed reporting to senior management. Such mismatch could have been problematic in other circumstances and it may have been that a more complete methodology, such as KADS with its Organizational Model, could have anticipated such a situation.

Management Lack of Recognition or Esteem of Expert

This situation was the key factor which prevented the Engineering Preventative Maintenance investigation and implementation proceeding. However, but to a lesser extent, it was also present in Benchtop Electronic Diagnosis. It is suggested that this problem is to be found in many "technician expert" cases where the professional engineers who are managers cannot believe that their technicians actually have the key knowledge required for an expert system.

(Alternatively, and taking a more sinister explanation, it may be that in some instances management are not prepared to accept the significance of the knowledge held by their technical staff since this could jeopardise the power relationships in the organisation.)

In several cases the author has had to politely excuse himself from polite discussions in manager's offices so that he could undertake the acquisition of knowledge on the shop floor. Although he failed in Engineering Preventative Maintenance the author suggests that this problem can usually be overcome once recognised.

Management Self-Satisfaction

In Telecom Repair Service the management seemed to be so pleased with their other efforts that they could not be persuaded to realise the extra advantages that could be offered by an appropriate expert system. Possibly this is related to the "Lack of Esteem" issue discussed immediately previously where engineering managers regularly do not really seem to appreciate the ability and knowledge of their technical staff.

Difficulties in Proving System Viability

Commercial expert systems implementations need to be financially sponsored by the management of the department where the expert works. In many cases it is very difficult to obtain access to sufficient expert time (and/or funding for the knowledge engineer's time) to undertake a feasibility or scoping exercise in order to sell the system concept to the sponsor.

Whereas a sponsor will readily commit funds in advance for a billing system or a database project, most potential sponsors remain very dubious about the merits of expert systems and will not provide funding unless there is some advance "proof" of the utility of the system. One of the reasons why the prototyping approach, with all its defects, has taken hold in the world of expert systems is that a prototype or concept demonstrator system is often necessary to give the sponsor some impression of the final system.

This is a very serious problem in the commercial world of expert systems. Unless the system is simple enough or there is sufficient access to the expert it is impossible to provide sufficient evidence to get a sale. (A "sale" may be to an external client or to another internal funding department.)

Access to the expert in advance of the funding decision is usually a key problem and the author has rarely had more than a day's expert access in order to acquire sufficient knowledge to either build a concept demonstrator system or make a firm costed proposal to build all or a major part of a target expert system.

In Chapter 4 we describe techniques for Exploratory Knowledge Acquisition that were by force adopted in order to overcome the problems described above. Had the author not adopted techniques to minimise expert time commitment in the initial stages of knowledge acquisition then it is likely that many of the investigations reported in Chapter 3, and the consequent implementations where these took place, would not have been undertaken.

6.4   ISSUES OF ECONOMIC VIABILITY

An expert system may be technically feasible and the potential sponsor may fully realise its utility and worthwhileness. However, if it does not "cost in", or the sponsor is not a funding government department, then it will not be built.

We consider issues of expert systems economic viability under the separate headings of "Construction" and "Maintenance".

6.5   ECONOMIC VIABILITY - CONSTRUCTION

Overall Cost-Effectiveness Compared to Unsupported Worker

In this case the expert system is either inherently too expensive to implement and/or would not be cost-effective compared to unsupported problem solving. This is probably the largest category of feasible systems never even considered for implementation. The author has on several occasions had to produce figures demonstrating that a lower grade of user (eg a "semi-expert" as described later in this chapter) can undertake a higher grade of work more quickly and/or with better overall results using an expert system and that the economies or increases in performance outweigh the direct and associated costs (eg hardware) of the system. Only when such cost-benefit type analysis has been completed can a successful case be made to fund the implementation. Even where the lifetime costs of an implementation are justifiable there is often a barrier in getting the capital expenditure costs needed to build the system. Many companies just do not have the often very considerable "up-front" funds to invest in

improving their performance with the aid of expert systems.

As discussed below, cheap delivery hardware is solving part of the problem but the knowledge acquisition "barrier", where skilled staff are required to occupy considerable amounts of their own and the expert's time, is still a key issue. As discussed elsewhere the author hopes that KADS and similar initiatives will be successfully adopted by practising knowledge engineers. However, such methods and the toolkits that go with them should perhaps be designed to offer rapid efficient implementations of straightforward expert systems rather than elegant implementations of complex systems. KADS in its elaborated form seems rather time-consuming and commercial colleagues of the author who are skilled KADS users (and ESPRIT partners) will admit that they cannot afford to use the full methodology but concentrate primarily on the Model of Expertise.

Modifiable generic expert systems (like expert 4GLs), rule induction and DIY systems (ie systems constructed and modified by the expert) may eventually provide an economical answer for a proportion of potential systems and may increase the cost-effectiveness of implementations in due course. However, for the time being and for the majority of tasks, knowledge engineers command good salaries and time to implementation is often quite protracted. Unless these problems can be overcome many potentially useful systems will be rejected by management on quite legitimate cost-effectiveness grounds.

Hardware Cost Issues

All expert systems require hardware to run on and if the user does not have regular access to a computer then the additional cost of providing

such hardware can prevent a system being implemented. In some cases such issues can be circumvented and the Audit Work Planning system was authorised, and was a considerable success, simply because it could be delivered "hardware cost free" on an existing PC available to the audit manager.

However, in the case of Benchtop Electronic Diagnosis there was great difficulty in gaining access to hardware in the workshop environment. Further, management were extremely reluctant to purchase benchtop PCs for each technician when they perceived that the implemented systems would be used as "Tutorial Expert Systems" where users only accessed the system on increasingly infrequent occasions when a fault to be diagnosed was outside the knowledge already learned from the expert system. The largely unexploited use of dedicated expert systems for training is considered below.

Expert Systems for Training - "Context-Sensitive CBT"

As described above, an expert system whose purpose recedes largely to that of a training system may not win support from a commercial sponsor whose initial declared interest was in an expert system to support less-experienced staff. However, such sponsors, usually via their training departments, are often keen users of Computer Based Training systems (CBT). In terms of overall business economy and training efficiency it might be sensible to evaluate the potential of expert systems which have an acknowledged training purpose in addition to their traditional advisory role. If this had happened with Benchtop Electronic Diagnosis, as described above, then local engineering managers might have benefitted from an internal subsidy from

the training department to lessen the unwelcome burden of additional hardware costs.

In addition from the potential economic benefits of using expert systems for both advisory and training purposes, there may be worthwhile training advantages from delivering training "at the point of need" rather than in a separately organised CBT or traditional training course. This topic is of potential interest for future research and the following suggestions are based purely on the author's personal experience as a trainee on CBT and traditional occupational training courses and his observation of the training users of expert systems for tasks such as Benchtop Electronic Diagnosis.

In the traditional "master and apprentice" form of training the "apprentice" receives some initial tuition and is then allowed to undertake some of the more simple or contributory tasks for his "master". From then on he usually interrupts his "master" when he needs additional information or has reached an impasse in the work he is undertaking. This process continues, with many interruptions in the early stages but later tailing off to the occasional request on a more complex and obscure case. The "interruptions" just described are training "at the point of need" which contrasts to CBT or classroom training. In the latter, training is necessarily delivered "in bulk" and in a context not relevant to the trainee's current need. In contrast, the "master and apprentice" form of training delivers appropriate small fragments of training in a context where the trainee needs to use the imparted information to undertake his immediate work. Intuitively, it would appear that the latter process is more likely to lead to better retention of the new material than might be obtained with CBT or classroom training.

In the same way, it could be argued that expert systems already offer

346

training in the same way as the traditional "master and apprentice" style of training. (This was the key "criticism" of the Benchtop Electronic Diagnosis system.) In other words, the key factors are that the training is delivered on a fragmentary basis at the point of need over a long period concurrent with the trainee's actual work.

In connection with commercial work outside that reported in this thesis the author has attempted to promote the idea of "context-sensitive CBT" as an extension to traditional expert systems. Such ideas have not received funding and no practical progress has been made. However, the key ideas are simple. A user would have access to an expert system to which he would seek recourse when he needed advice on how to tackle an unfamiliar problem or when he reached some other form of impasse in his immediate work. Initially the system would follow the traditional expert system dialogue to the point where the immediate problem had been overcome. At this stage, the system would offer the user the opportunity of a short CBT session covering the wider principles of the case just covered. Alternatively, the user could opt for this session to be added to a training agenda of topics which would be typically covered in a short worksite training session at the end of the working day.

The implementation of a context-sensitive CBT system linked to an advisory expert system (or their complete integration as an extension of the expert system's "Help" facility) would seem to be relatively straightforward and could radically improve the acceptance of both expert systems and CBT in the workplace. The impediments would seem to be mainly organisation with the possibility of powerful training lobbies in commercial organisations unwilling to give up their rights to withdraw staff from their daily duties and provide training services for them.

Cost-Effectiveness of "Partial" Expert Systems

As discussed in Chapter 3 a Job Shop Scheduling expert system would have required a human "scheduler's assistant" to telephone customers to negotiate changes in delivery dates in order to avoid "bottlenecks". In management's view if a skilled person is required to service the scheduling system then that person can actually undertake the scheduling process even though the expert systems scheduler would actually produce better, faster schedules.

Thus if the expert system cannot support the "whole task" then it may not be economically viable to implement.

6.6   ECONOMIC VIABILITY – MAINTENANCE

Maintenance of Changing Knowledgebase

This issue has already been raised under "Lack of Stable Expertise" in respect of Eurobond Warrant Trading and Benchtop Electronic Diagnosis above. Although the issues raise by Eurobond Warrant Trading seem impossible to reconcile, the type of problem demonstrated by Benchtop Electronic Diagnosis is very common and is discussed further under "Opportunities for Communicative Expert Systems" towards the end of this chapter.

Localisation Issues

The successful deployment of an expert system can often involve

348

maintenance modifications in the form of "localisation" to the needs of geographically remote parts of the organisation. For example, in tasks not reported separately in Chapter 3, a supermarket shelf layout ("facing") system had to allocate product lines optimally for different sizes and layouts of supermarkets. Therefore the system required a version for every store embodying the data on the store shelf dimensions and local sales priorities so that correct layouts could be produced. The effort to produce such versions is very considerable.

Similarly, in a system that was investigated but never implemented, a well-known high street store wanted to allocate display areas for product lines according to fashion season. For instance, this involved increased space allocation for confectionery at Easter, more space for toys at Christmas and more space for holiday clothes in the Spring. Such "seasonal localisation" would have required regular maintenance updating changes to the system's knowledge base and this was a prime reason why the system was declined on economic grounds.

In Business Performance Assessment the knowledge base had to reflect different evaluation priorities for different health areas. For instance a greater proportion of geriatric patients was acceptable in the South Coast retirement area than would be allowed in a Northern industrial town. Similarly performance targets required regular maintenance updating so that different specialisations could accept, for instance, different times to discharge. Thus the post-surgical target discharge period is longer for orthopaedic patients than gynaecological patients and the ratio between the two is constantly changing with advances in care. Such changes impose a system maintenance load which, in the case of Business Performance Assessment, was readily accepted by the sponsor.

6.7   SPECIFIC KNOWLEDGE ACQUISITION ISSUES

Lack of Expert Cooperation and/or Access

This is a straightforward problem and the reasons have been largely covered above. If the expert cannot be accessed or is uncooperative then no system can be built. The author has met very few uncooperative experts and, where this has been an initial problem, it seems to have usually resulted from factors such as job insecurity on the part of the expert.

(An example of one such rare case was met in an electronics repair workshop in the Midlands, not related to any of the tasks reported in this thesis. On his first visit to undertake knowledge acquisition, the author was greeted by the local union representative who was concerned that the proposed expert systems work was connected with threatened regional redundancies. Once the necessary, and very truthful, assurance was provided the expert and his union colleague proved very cooperative and genuinely interested in the advantage that expert systems might provide for the workshop staff, perhaps as an aid to enabling them to keep their jobs.)

Lack of Consensus on Expertise

Disagreement between the experts was a major problem in Pensions Management where interpretation of Pension Scheme rules differed between experts from different parts of organisation. This issue was eventually solved by demanding that the sponsor appointed a single expert to coordinate the views of the many experts.

Acquisition of "Perceptual Knowledge"

In Chapter 4, in the context of Eurobond Warrant Trading screen interpretation, we discuss the difficulties of acquiring "perceptual knowledge", which may be equivalent to the Dreyfus and Dreyfus (1986) term "arrational knowledge". We have not investigated a similar case except for a brief investigation of the skills of a sales representative who had an excellent reputation for developing sales to the shops associated with petrol stations. In his case he constantly referred to "recognising the clues in the customer's face and posture" during a sales visit but could only give the barest clues to what he actually meant.

The investigation of perceptual/arrational knowledge is outside this research. However, where a key part of the expert's performance depends on such knowledge then it may be difficult to acquire it for the purposes of building an effective expert system.

6.8   SPECIFIC USER INTERFACE ISSUES

The user interface is the only part of an expert system seen by the user and unless it is appropriate and acceptable then the system will fail. User interface issues for expert systems has been the subject of extensive comment and research, eg Carroll & McKendree (1987), Norman (1983), and extends the more general work on user requirements for more conventional systems, eg Williges & Williges (1984), and the specialist work to extend expert system and other interfaces such as the work on hypermedia based on Nelson (1965) and extended as in Nelson (1973).

Protracted or Inconvenient Dialogue

If the dialogue with the expert system is more "time expensive" than unsupported solving problem then the system will fall into disuse. For instance, in Benchtop Electronic Diagnosis the lack of a convenient benchtop expert system caused the user technician to prefer to randomly substitute groups of components rather than start a dialogue on the workshop PC. However well a user interface is designed, some time and effort is required for user interaction with an expert system and a user may perceive that this is longer or more inconvenient than undertaking the task itself, even if this latter is not very efficient.

Similarly, in an investigation made by the author reported elsewhere in this thesis, a company credit rating system was not implemented since its usage would have required users to type into the system the handwritten responses on enquiry forms from companies seeking credit before ratings could be computed by the expert system. In the absence of a readily available and effective handwriting recognition system, such effort would have been unacceptable to users and the system was not built.

User Engagement Issues

Unless it is a compulsory aid to a user's work, an expert system will not retain its users if it does not sufficiently engage them to continue with the dialogue. Such potential problems can be overcome and a good example is provided by the Small Business Advisor expert system which was successful because it engaged the "short attention span" businessman user by early partial feedback of results.

Support for Existing or Natural Working Style

Transfer of users to work supported by an expert system will not be straightforward if the system does not support the users' natural style of working or normal user documentation. Thus the Commercial Business Underwriting system depended for its success on capturing the "virtual form" always used by expert. Similarly in Multinational Tax Planning it was necessary to show and print familiar balance sheets and schedules to the user expert. In the case of scheduling system it is necessary to use the familiar Gantt chart for communicating with the users since this is the universal way of representing job shop schedules.

For similar reasons a system will probably be unacceptable if the output or printout is in unfamiliar format or not of sufficient print quality. Thus the Audit Work Planning system was successful because it entirely replicated "manually produced" output including the cosmetic details of the forms used by the audit manager.

Again, and somewhat perversely, management required us to simulate the low quality mainframe "green screens" in the Life Underwriting expert systems so that its appearance would be compatible with the rest of the mainframe application processing system.

Verbal/Visual Input Encoding Issues

A system which depends on the laborious textual encoding of verbal or visual information seems unlikely to be acceptable to users. This has already been covered in the discussion on Eurobond Warrant trading where

remarks on prices from counterparties would have had to be fully encoded as text.

User Privacy and Convenience

All expert systems need to be convenient to use as with Audit Work Planning which had to be portable for use at clients' premises. However, some financial systems require the input of considerable private information and, for instance, it was a key requirement of the Small Business Guidance system that it could be used in the privacy of the small businessman's home or office rather than at the government agency's premises. This in turn had implications for producing the system so that it would run on a cheap PC such as that possessed by a small firm.

6.9   SPECIFIC ISSUES OF EXPERT-USER APPROPRIATENESS

Expert systems must satisfy and match the needs of a range of "experts" and "users" if they are to be usefully accepted. Previous research on "experts" and "novices" has been driven primarily to discover the differences in the knowledge representations or mental models used by, for instance, individuals with different levels of understanding of domains such as physics (Chi, Feltovitch & Glaser, 1980), chess (Chase & Simon, 1973), financial problem solving (Hershey, Walsh, Read & Chulef, 1990) or computer programming (Kahney & Eisenstadt, 1982; Soloway, Ehrlich, Bonar & Greenspan, 1982).

In this work we discuss our findings in respect of expert and user

354

categories with the rather different motivation to determine what expert system support, if any, might be appropriate for such categories. However, as a short background preface, it is useful to briefly recount how this particular topic was important at the commencement of this research.

Background – "Star" Diagnosticians

As explained in Chapters 1 and 3, this research originated in the study of engineering problem solvers undertaking diagnostic tasks for British Telecom. Since then the scope and objectives of the research have very much broadened.

Nevertheless, probably the initial focus of interest (and one that, regretfully, has not really been fully followed up) was to understand the differences between "stars" and ordinary diagnosticians.

Initially in the RSCs (Repair Service Controls) where we investigated the Telecom Repair Service task and, later, in the regional workshops where we investigated Benchtop Electronic Diagnosis, managers would often use the term "stars" when describing their staff and suggesting those who might be used as the subjects for knowledge acquisition.

In a workshop of perhaps ten diagnosticians of the same rank, it usually seemed that one or two would be categorised as "stars". The same sort of proportion applied in RSCs where perhaps three or four staff would be identified as "stars".

"Stars" would not only be identified by their bosses but very readily by

their peers. These were the people that everybody turned to when confounded by a new or difficult case. What made them different was an unfulfilled early objective of this research. Maybe they were more intelligent. Or perhaps more genuinely fascinated with their work. Because they were the subjects for our knowledge acquisition we cam to know these individuals very well and to respect them considerably. Although we never properly investigated what, if anything, set them apart from their colleagues, it seemed that they did have three superficial characteristics in common. First, they were usually quiet and modest about their abilities. Second, they were inclined to take notes and were tidy and neat in their reports. Third, they seemed genuinely interested in their work and had a very good and detailed understanding, albeit at a "craft" rather than "theoretical" level of their subject.

The term "star" or anything equivalent to it, has never been used by managers or staff in the investigations of non-engineering tasks. Nevertheless, we have recognised individuals (often, again, those who were recommended as subjects for knowledge acquisition) who would fit this term. Thus, in particular, Audit Work Planning and Pensions Management were tasks where "stars" could be recognised in much the same way as for engineering tasks.

The issue of what differentiates "stars" and "others" has not been properly pursued in this research and is not dealt with further. However, it has been increasingly clear that the simple words "expert" and "user" are woefully inadequate to describe the many types of persons falling into these categories and thus, at a very much more simplistic level than might have otherwise been the case, this issue is covered in the remainder of this section.

"Experts"

We suggest that the term "expert" has some limitations in respect of its usage to describe the several types of person who might come into contact with an expert system or be considered as a source for the knowledge contained within it. In particular, we suggest that the simple term "expert" does not adequately quantify the "creativity" dimension of expert performance.

In its common occupational usage, the word "expert" is used to describe individuals with rather a wide range of abilities. At one extreme, and where the word "specialist" is sometimes used and is actually more appropriate, some "experts" are just the result of training courses and direct tuition from colleagues (eg apprenticeship). Although their performance may be as good as any other "expert" they may not have the aptitude or ability to regularly and creatively extend their knowledge. We have termed such experts "competence experts".

In contrast, and at the other extreme, there are those "experts" which we term "knowledge creator experts" who, whether from first principles or by "trial and error", are the persons who "break open" a new task and establish it on a routine basis.

The particular, term "break open" has been used by managers to the author on several cases, particularly for engineering tasks, to describe the way in which they use what we refer to as knowledge creator experts to introduce a new task into the organisation's repertoire. Benchtop Electronic Diagnosis and Eurobond Warrant Trading were both cases where management assigned particularly able experts to new tasks.

Once the tasks were established, lesser experts (which we term "competence experts" or "semi-experts" below) would undertake the tasks so that the best experts could move on to "break open" yet other new tasks. For instance, in Benchtop Electronic Diagnosis, two particular technicians were regularly employed (and rewarded) to undertake the repair of new items coming into the workshop and, once a method of diagnosis had been established, to hand over the task to less experienced colleagues. Some of these latter technicians were either learning the job or, because of their limited abilities, were only given the simplest cases with which to deal.

Similarly, in the case of Eurobond Warrant Trading, the expert assigned to us was very much a specialist at "breaking open" the trading of new combinations of instruments. In the competitive banking situation it is necessary to entice customers to invest with an ever-evolving range of financial products. The bank must both promote these new products and, in parallel, must learn how to trade the underlying instruments properly. As part of the preparatory briefing for the knowledge acquisition for Eurobond Warrant Trading the expert told us that he had been assigned the task of establishing a series of new products and for "looking out for" opportunities of trading that would lead to new and competitive products. Once trading was relatively established he would gradually transfer the new instrument to colleagues and move on to new ground.

Again, in Pensions Management a few of the regions were served by Pensions Officers with very little experience or aptitude and it was their practice to refer to their headquarters colleagues anything except the most straightforward enquiries. In this case the headquarters Pensions Officers not only dealt with the more complex routine enquiries but also dealt with novel pensions problems from the whole country. In many cases such novel

cases lead to an extension of the corporate "pensions scheme case law" and regular updates were sent out to the regional Pensions Officers so that they could deal with such new cases if they encountered them. However, it remained only for a few headquarters Pensions Officers, and a handful of very able regional colleagues, to create new pensions scheme knowledge in this way.

There is it seems a very real confusion between the term "expert" as used in this thesis (ie the holder of valued knowledge) and the use as automatically applied to a job grade. Just because individuals may have the title "Pensions Officer" or "Audit Manager" does not necessarily mean that they are "experts" (as we use the term) in these fields. However since their pay and job grading depends on the assumed possession of "expertise" they are regularly referred to as "experts" by their colleagues and employers.

"Users"

This thesis and the commercial work related to it have been largely concerned with experts and expert systems. However, particularly in the knowledge acquisition and commercial implementations, we have had many opportunities to understand the social and organisational setup of individual offices, workshops and other workplaces and to meet a considerable range of potential and actual expert systems users. This section comprises a short overview based on that experience.

Just as we can differentiate categories of experts so to can we see "users" (as applied users of expert systems) split into a number of separate

categories. In some cases, "users" are fully-fledged "competence experts" and, as described above, have been taught their expertise but have no aptitude to expand it. In such latter cases they do not need expert systems and, usually, do not provide good sources for knowledge acquisition.

Some users are "semi-experts" who, although they have some worthwhile knowledge are not competent to reliably undertake the task. For instance, in a repair workshop setting, these might be technicians who are experienced at repairing other equipment but have only sparse experience of the equipment being made the subject of the expert system.

Finally, the lowest category of user is the "problem-naive" user who is totally reliant on direct instructions from an expert or, in a few cases, guidance from an expert system. Trainees and junior apprentices make up the majority of this category in the engineering context. In the non-engineering area such persons are usually from other departments to that where the expert task is undertaken. Foe instance, in Pensions Management, Personnel Department clerks would utilise the skills of Pensions Officers but had little understanding of even routine pensions scheme rules and regulations.

Of course, the categorisations of experts and users suggested above is rather imprecise and individuals do not always fit easily into a particular category. Similarly, individuals develop and change and particular individuals can change between categories. For instance, "semi-experts" or "competence experts" sometimes overcome gaps in their inherited knowledge and, in due course become sufficiently experienced to develop their own independent expertise. In these instances such they can be considered to have become "knowledge creator experts" even though, on the basis of their

prior history, they would not have seemed to have either the aptitude or motivation to reach this level.

A Hierarchy of Expertise

In summary we recognise "a hierarchy of expertise" embracing the above "expert" levels and also levels of those who are unequivocally "non-expert". Our categorisation is as shown below:

Knowledge Creator Experts who are both fully task competent and capable of extending and improving their existing expertise.

Competence Experts who are fully task competent but have little role or aptitude for creating or developing new expertise

Semi-Experts who have some task expertise which may however be incomplete or unreliable

Non-Expert Workers who know or are associated with the task but have little or no worthwhile task expertise

Problem-Naive Persons who have no knowledge or experience of the task

(We have deliberately avoided the "Expertise-Proficiency-Competence-Advanced Beginner-Novice" categorisation proposed by Dreyfus and Dreyfus (1986) as we feel that the more straightforward definitions above are very

sufficient for our purposes.)

The Needs of Occupational Experts

We discuss the needs of the various categories of users of task expertise in the following paragraphs.

Expert Systems for Experts?

All of our commercial expert systems implementations, including those progressing from the work reported in this thesis, have been primarily designed to allow persons other than the existing expert to undertake the designated task. The availability and use of an expert's knowledge by another person separated in time or space from the original expert seems almost to be part of the implicit definition of an expert system. In fact many of the claimed commercial benefits of expert systems, such as the use of a less experienced (and therefore less expensive) person to undertake a task or the ability to undertake the task at multiple or remote field locations, assume a difference in identity between expert and user.

So what, if anything, is the role of expert systems for existing experts in the same field, including the expert or experts from whom the knowledge was originally captured? In the case of the more straightforward tasks it would seem that experts rarely need expert systems. For instance, the expert for Benchtop Electronic Diagnosis, although a most kind and knowledgeable man, was barely interested in the expert system based on his knowledge. He would admit the system could be useful to others but could

see no use for it for himself. Similarly, the sponsoring partner and manager for Audit Work Planning were enthusiastic for the roll-out of the system to their colleagues but never really considered using the system themselves. In general, for such tasks, experts will usually admit that expert systems are "clever", "interesting" or "useful", they often cannot see how they could perform any practical purpose in their own day-to-day work.

The situation was similar with some of the more complex tasks such as Pensions Management and Life Underwriting. In such cases the key experts were totally in command of their tasks and saw little or no reason to use the systems which they were so keen to see used by their colleagues. However, this finding was not universally true and the experts for Small Business Guidance and Multinational Tax Planning were keen to be users of the systems which they themselves had helped to create.

In the case of Small Business Guidance the final system was based on the knowledge of four primary experts each with rather different, but overlapping, fields of expertise in the area of small business analysis and advice. (In Section 3.11 we refer to the "theoretician", "counselling expert" and the "generalists".) These experts saw the value to themselves of the final system and, although not showing overwhelming desire, were reasonably keen to experiment with its use in actual client situations.

For Multinational Tax Planning there was an even greater interest to use the system incorporating the knowledge of several contributing experts. As explained in Section 3.7 the system could be used to give tax advice for a different continent to that with which the user expert had familiarity. Therefore the experts in this case could very much see the possibility of

making unaided multinational tax optimisations involving, say, American and Australasian subsidiaries where their own knowledge was more specific to Europe.

Further, and for a reason different to Small Business Guidance, the sheer complexity of some of the multinational tax problems (and the calculations associated with them) could be more efficiently handled by the Multinational Tax Planning system event though they were in the area of primary competence of the user expert (eg restricted to North American subsidiaries) and would have normally been dealt with on a laborious basis by the expert.

Further, and again as exemplified by Multinational Tax Planning, the excellence of the user interface in the final commercial system made it attractive to use by experts since, with great ease, they could experiment with alternative corporate holding structures to determine that which was most tax efficient.

In summary, and from the experience of this research, it is justifiably difficult for experts to see how straightforward expert systems based on their own knowledge can assist them. Thus, for knowledge creation experts in tasks similar to Benchtop Electronic Diagnosis, the experts are custodians of existing expertise (often of their own origin), have the ability to extend it and fail to see how expert systems technology, which so far has just tracked their historical knowledge, can support them. However, when the expert system has some added value, perhaps by incorporating the knowledge of peer group experts or by handling or making explicit particularly complex problems, experts will often wish to experiment or use the system to which they have contributed.

Existing expert systems technology has largely addressed the challenge of enabling less-experienced or unexperienced users to employ the knowledge of experts. In the main, the technology has only indirectly approached the needs of knowledge creators or competence experts and there is some considerable scope for research in devising systems to support such experts.

The "Hard-Wiring" of Existing Knowledge

As discussed above, systems such as that implemented for Multinational Tax Planning enable both users and experts to explore the solutions to novel problems. This is in contrast to the author's experience of many commercial expert systems which do little more than connect a pre-determined set of inputs and outputs and so have no chance of providing new solutions to novel problems.

In many such commercial systems, including it must be admitted several of those where the author has had an involvement, the content of both the sets of input and output conditions, and the intermediate linkages, have been fully "designed" and tested by the knowledge engineer. In most cases the knowledge is purely specific and associational with little or no causal knowledge to allow any extension of the inferential possibilities envisaged by the system's implementor.

The "hard wiring" of inference control through the knowledge base is also common in commercial expert system implementations. It is no wonder that so many knowledge engineers report that experts react to the viewing of "their" expert system with the response "I could have written it in BASIC".

Lack of World Knowledge

In many cases the knowledge creation limitation results from the practical impossibility of encoding the world knowledge available to the human expert. For instance, in Permanent Health Insurance, it was only the human expert who could realise that a family history of epilepsy had potentially disastrous consequences on the career of a person who gave their occupation as "electrical power engineer".

In this actual case the human expert was able to predict that a person whose work involved close proximity to high voltages would be unemployable if they developed epilepsy. No realistic expert system would have had sufficient world knowledge to make this inference and, in consequence, set an increased premium.

Lack of world knowledge is a serious practical limitation which is unlikely to be overcome in practical expert systems in the foreseeable future and, probably more than any other reason, reinforces the need to retain human experts in their creative role, either where that creativeness involves new knowledge or, with some important exceptions, novel solutions. Such human experts may, as with Multinational Tax Planning, work with expert systems but the end result will be an improvement of problem solving efficiency or quality rather than the replacement of the expert. The particular issue of problem solving quality is discussed below.

Expert Systems for Maintenance of Problem Solving Quality

By their nature, knowledge creator experts are usually vigilant to less

common conditions and respond accordingly. In fact their creativity usually relies on such vigilance.

Nevertheless, even experts of the highest calibre seem to show variations in problem solving quality. For instance Telecom Repair Service in particular showed us that experts of apparent equal knowledge showed considerable differences in problem solving quality. Thus, based on the same experience, some RCOs were far more likely to make novel and correct diagnosis than other RCOs who, apparently, took less trouble and gave "default" standard diagnoses which did not take into account the novel features (or clues) in the problem.

In fact during our initial observations, we noticed an apparent tradeoff between problem solving quality performance and problem solving rate. (This, of course, was also seen with Life Underwriting where management were firmly in favour of a high turnover even if this lowered "quality".) Thus, in Telecom Repair Service, those RCOs who achieved the most insightful and correct diagnoses usually made less diagnoses in a working day. Nevertheless, if challenged, the faster and apparently less able diagnosticians were able to make equally creative and insightful diagnoses as those RCOs who were slower and, apparently, more creative in their problem solving. Such differences did not seem motivated by any perceived pressures by management to either offer "quality" or "quantity". In fact the differences were first pointed out to the author by the RSC engineering manager who expressed only interest and certainly no negative opinions to either group of engineers.

As a result of these semi-formal observations, and similar observations from some of the other tasks investigated, we suggest that there is a need for

367

systems to support problem solving quality. In particular there seems to be a tendency for some experts to ignore clues and move directly to a quick "convenient" straightforward solution.

Expert "Sub-Systems"

Part of the reason why an expert's quality may drop is that the expert has to cope with the problem solving work associated with his existing knowledge and, at the same time, cope with parallel inferences where the knowledge is not so well established. Thus there may be a real opportunity for "expert sub-systems" to which creative experts can "devolve" some of their more mundane problem solving activity.

Taking this idea further, and envisaging such systems in place, they would have a useful role in prompting the expert to partial problem conclusions. Thus, in a manner analogous to a blackboard system, semi-independent knowledge sub-systems (relying on established "knowledge sources") would operate in the background and throw partial problem conclusions into the human expert's workspace to enable him or her to make the final creative (or non-creative) leap to a problem solution.

The Needs of Knowledge Users

So far we have we have discussed the needs of knowledge creators. In the following paragraphs we consider lesser "user" categories of knowledge workers.

Quality Control and "Checklisting"

We have already discussed above some of the special problems of "Maintenance of Problem Solving Quality" with special reference to knowledge creator experts.

The maintenance of problem solving quality perhaps applies even more acutely to the "semi-experts" who support and learn from the knowledge creators and other fully experienced experts.

In Benchtop Electronic Diagnosis in particular there was evidence of considerable lack of care and rigour on the part of some "experts". In this and other cases management often attempt to try to introduce checklists and other quality control procedures in an attempt to force staff into rigorous practices. This has already been discussed separately in Chapter 4 on problem solving props. The use of such checklists and other devices points to a clear need for support to problem solving rigour in the knowledge user community.

In some cases, even though some of their "expertise" is little more than checklisting, expert systems (where the checklists are essentially hidden) provide a far more socially acceptable means of maintaining problem solving quality than more traditional checklists which are often resented and poorly used by workers.

Support for Problem Solving Props and Coping Strategies

Checklists, which are mentioned above, are a useful form of problem solving

369

prop. As discussed in Chapter 4 there are a range of existing problem solving props which either experts have devised themselves or are institutionally provided to facilitate their work.

In recognising that experts have an existing need to use such props we suggest that such needs might be better met as part of an expert systems user support system. Such support will probably be in a slightly different or "intelligent" form to those used by unsupported experts.

Flexibility of Usage

In some tasks, such as Benchtop Electronic Diagnosis, Audit Work Planning and Life Underwriting, the expert is free to adopt a pre-ordained "procedure" to problem solution.

Thus in Benchtop Electronic Repair the diagnostician is simply faced with a non-operative unit and can set about diagnosis in a routine and regular pattern. As a result, the TRACKER expert system could offer such advice to a user with no requirement for flexibility of problem solving procedure. Similarly an audit can, and is, planned in an orderly and methodical manner on the basis of pre-existing facts and Audit Managers will tend to solve the "audit planning problem" in the same way on each occasion. The underwriter's problem solving in Life Underwriting is even more routine.

As a result the expert systems that support these tasks need only embody a single route to the same solution. In contrast other tasks ideally demand more flexibility in usage.

370

For instance in Pensions Management the Pensions or Personnel Officer has to answer the Scheme Member's enquiry based on details sometimes provided verbally and sometimes in writing. In such cases the human expert is able to flexibly alter his or her problem solving to take account of the different ways in which information is provided and requested. The implementation for the Pensions Management expert system offers some flexibility in this respect but it is not a feature regularly found in expert systems and, even in this case, the flexibility does not nearly approach that provided by the human pensions expert.

Recognition of "Good Solutions"

In the investigation of Job Shop Scheduling we recognised that expert schedulers were far more competent at recognising a "good schedule" than generating one. Thus, although schedulers may have little ability to find a good solution unaided, they are usually able to say whether a schedule is "a good schedule".

A simple analogy would be to suggest that the situation is similar to the individual who cannot solve a cross-word puzzle yet is easily able to recognise and appreciate a good solution. In the case of the cross-word puzzle a "good solution" would, for instance, necessitate that words interlock and are valid solutions to the clues.

However, in the case of schedulers the recognition of "goodness" of the solution does not merely depend on the simple mechanical checking that the schedule meets the various constraints etc. There seem to be some "aesthetic" features of a good schedule which schedulers recognise over and

above its simple solution to the initial problem.

In the investigation of Job Shop Scheduling, scheduling experts did not seem able to express their knowledge on what constitutes a "good schedule" over and above the simple meeting of constraints. However, their ability to recognise good schedules, even though they could only partially explain what made them "good", seems a useful form of knowledge which could be exploited to evaluate schedules produced either manually or by a knowledge-based scheduler.

The partial explanations of "aesthetic goodness" pointed at issues such as the load balance on individual resources, the apparent robustness of the schedule to the need to accept urgent orders and cope with resource breakdowns, and the extent to which the schedule meets due dates "just in time" rather than completing orders well ahead of requirement.

Obviously such factors represent additional constraints, over and above straightforward meeting of due dates etc, which might be incorporated in an expert scheduling system to produce schedules which not only satisfied the explicit constraints but also the "good schedule" constraints of the expert scheduler. Knowledge relevant to the "good schedule" constraints might be inductively acquired by asking expert schedulers to score schedules and, once made explicit, could be incorporated into the expert system.

Summary

We reject the simplistic idea that "experts" are the source of knowledge for

expert systems used by "users". We suggest that there are a range of categories of experts and users to be considered and that, although it is still difficult to properly support "experts", this and the issue of support for different types of users is an essential step in ensuring the appropriateness of an expert system for its environment.

6.10 SPECIFIC HARDWARE ISSUES

We have already discussed the economics of hardware provision above. Nevertheless, apart from cost, hardware issues are often crucial in commercial expert systems work. For instance, in Telecom Repair Service the expert system could not be ported to the Honeywell hardware supporting the record keeping software used by the diagnostician. In addition to other problems this prevented the expert system being considered for full field usage.

In other cases the hardware issue has been solved. In Life Underwriting it was difficult and time-consuming to build a system in IBM's ESE and then Aion's ADS on the IBM mainframe computer. The project involved much lost time because ESE did not support re-entrant code or knowledgebases and therefore was useless for more than just a few users. Nevertheless, once this problem was solved by re-implementing in ADS, the expert system moved immediately into widescale productive usage with the company's many mainframe users.

Although not strictly a hardware matter, interfacing and integration issues can often take more time to complete than the expert system itself. This was particularly true for Life Underwriting which involved very

considerable usage of mainframe data. Similarly, interfacing to information feeds and sources would have entailed major work in Eurobond Warrant Trading.

6.10   OPPORTUNITIES FOR COMMUNICATIVE EXPERT SYSTEMS

The report on Benchtop Electronic Diagnosis in Chapter 3 discusses the topic of "Design-Diagnostic Expert Systems" to aid the communication between the designer and the repair technician who both have separate expertise concerning the same item of equipment. This example is but one of a number of opportunities for what we term "Communicative Expert Systems". As discussed above it is sometimes very difficult to prove the cost-effectiveness of a particular expert system especially when it does no more than support or replicate skills already existing in the business. Communicative Expert System, which follow the trend of distributed computing environment, would offer new functionality and could overcome one of the major obstacles to efficiency in commerce and industry. Many workers are established experts at their tasks but the institutionalised barriers to communication between different organisational hierarchies and different status groups prevents the useful flow of knowledge and information between workers who are very often separately dealing with different aspects of the same problem. The examples used in what follows come from the investigations reported in Chapter 3 but have been supplemented with other cases from common experience.

Design-Diagnostic Cooperative Expert Systems

This has already been fully discussed in respect of Benchtop Electronics Diagnosis. However, a similar situation must exist between the designers and repairers of domestic electrical/electronic equipment, motor cars and similar goods. For instance, repairers of domestic washing machines will always comment on the stupidity of certain design features, often involving the trapping and damage to wires and gaskets. Although Benchtop Electronic Diagnosis presented an ideal case for designer-technician intercommunication it seems likely that similar cases must exist where the ideas outlined in Chapter 3 could be applied.

Problem Referral Expert Systems

A need for a problem referral system occurred during the implementation of the Small Business Guidance expert system. Different experts are able to deal with special areas of a problem and, for instance, a generalist at business guidance would need to refer to a specialist for advice on a hotel business or a dairy farm. An expert systems implementation can, of course, cover the knowledge of all such "sector/niche experts" but this is expensive. In some cases what is required is an expert system to prepare a report covering the essentials of the problem to the sector expert so that appropriate guidance can be given on the special area.

A similar example exists between GP and Consultant where the GP prepares the information for the Consultant to make the specific tests and diagnosis. Currently this is usually accomplished via a letter, sometimes very brief and occasionally near-illegible. A better solution might be to have a case

referral expert system which prompted the GP to enter the key referral information and then which, at some later stage, prompted a nurse or assistant to enter more comprehensive relevant case details in order that the expert system could prepare a comprehensive referral letter to the Consultant covering the specific facts needed by the Consultant to complete the further diagnosis and treatment.

Of course, a problem referral expert system would only be possible if the generalist, GP or otherwise, were able to sufficiently recognise the broad symptoms in order to make a referral to the correct expert. In the case of Small Business Guidance we were confident that such referrals were competently made. However, in the main, the knowledge only involved recognising the business sector in which the businessman worked and selecting an appropriate sector expert with specialist knowledge of that sector. The referral task for a GP is often much more difficult and it would form an interesting area of research to determine the type of knowledge required for referral and how this differs from the knowledge to undertake the in-depth task itself.

Hierarchical Expert Systems

The term "Hierarchical Expert Systems" is used to refer to a special form of communicative expert systems replacing a series of hierarchical experts. An excellent example is provided and discussed in the section on Job Shop Scheduling in Chapter 3. Here the existing task was supported by an hierarchy of scheduling experts with a time delay between each level of the hierarchy. Replacing the hierarchy by a single expert system would have prevented catastrophic manpower wastages when urgent re-scheduling was

necessary because the original schedule could not be completed.

6.11   SUMMARY

This content of this chapter directly arises from the investigations reported in Chapter 3.   The history of expert systems is littered with examples of systems that were "never quite" finished or brought into active successful use.   Usually the reasons are non-technical and sometimes comparatively trivial.   However, the world of expert systems has been blighted by such failures and it is hoped that, by structuring the experience gained from the investigations reported in Chapter 3, it may be possible for other workers to avoid some of the pitfalls.

Nevertheless, some obstacles will remain, particularly associated with the cost of implementing and maintaining systems.   Such costs are unlikely to be reduced since knowledge engineering effort is expensive, even if it is being gradually being enhanced by appropriate non-intrusive supportive tools and methods.   In such cases the expert systems implementor must look to the other side of the cost-benefits equation to determine what additional advantages can be offered over and above the mere replication of existing expertise.   Such additional benefits might enable the financial justification more expert systems to be deployed in organisations which up to now have not seriously adopted the technology.

CHAPTER 7 - PROSPECTS FOR EXPERT AND KNOWLEDGE-BASED SYSTEMS

7.1   INTRODUCTION

This chapter takes a realistic view of the prospects for knowledge-based systems and, in particular, expert systems.   The view has been developed partially as a result of the author's own work, as reported and discussed in this thesis, and partially as a result of his observations of the work of the expert systems community of which he is member.

The commercial promise of expert systems has not yet been delivered.   Yet, and probably like anyone else who has delivered commercially successful expert systems, the author remains convinced that the technology has the potential to offer very great business advantages.   So, for the author, there is no question as to whether expert systems can commercially survive and flourish.   The questions that concern him are how and when this may happen and, very particularly, what are the factors that are preventing expert systems fulfilling the role that has been long promised for them.

The author suggests that the investigations in this thesis, and the results drawn from them, offer some clues in answer to these questions.   There are no dramatic answers and the overriding observations are simple:

1. The technology should not be used for tasks to which it is unfitted.

2. The methods, especially KADS, must be appropriate, or be made appropriate, for use in developing the required applications.

3. Knowledge must be considered as a component of conventional systems reflecting business needs and not solely for incorporation in KBSs whose construction is often driven by available knowledge.

4. "Design" must take an increasingly major role in the development of systems incorporating knowledge.

5. Knowledge based systems should be targeted not only at tasks where experts are competent but increasingly toward tasks where experts can never be competent.

7. Unsatisfied business opportunities, especially those involving knowledge at the corporate rather than individual level, should be enthusiastically addressed.

The author is optimistic that such conclusions, if followed, can lead to the much wider uptake of expert and knowledge-based systems technology. Each of the above observations is separately discussed in the following sections. For convenience in this chapter, and except where explicitly stated, the term "expert systems" will be used in a general sense to cover both expert and knowledge-based systems.

7.2   AVOIDING FAILURE FROM APPLICATION/TECHNOLOGY MISMATCH

Chapter 6 covered a range of factors predisposing to the feasibility, success or failure of an expert system project. Like any other technology, expert systems technology cannot be expected to be useful for tasks to which it is unfitted. Therefore, and taking up some of the points from Chapter 6,

intending expert systems implementers should not address tasks where the existing expert knowledge is incomplete as with Permanent Health Insurance or unreliable or unstable as with Eurobond Warrant Trading. And it is unwise to attempt to support problem solving where, as in Eurobond Warrant Trading, there is a key requirement to encode a mass of verbal and visual data. Similarly, the reputations of both the technology and the implementors will be damaged if technically successful expert systems are promoted for tasks where they cannot be cost-benefit justified.

Other detailed issues are discussed in Chapter 6. This thesis has covered some indicators as to tasks for which there are technical reasons making it unlikely that an effective expert system can be implemented. We have also indicated some of the organisational barriers that may prevent a potentially technically adequate system being implemented or adopted. What is not covered, and is not appropriate for this thesis, are the straightforward cost-benefit analysis methods which can very quickly indicate whether any proposed system has a chance of commercial success. Unless any proposed system can satisfy all of the technical, organisational and cost-benefits requirements then, unless a genuine research objective is being followed, it is foolhardy for the commercial implementor to proceed.

7.3   KADS - A CONSTRUCTIVE COMMENTARY FOLLOWING THIS RESEARCH

The KADS methodology, discussed in Chapter 2, was not used in support of the investigations reported in this thesis. However, it is useful to consider the findings of this research in relationship to the KADS methodology. In particular, what is the relevance of the problem solving components, progressively identified in Chapter 3 and summarised in Chapter 5, to the

knowledge used by KADS in the construction of a knowledge-based system?

At first sight, there would seem to be a very clear answer to this question. From the author's understanding of the KADS methodology it would seem that most of the components are really consequences of the cognitive limitations of the human problem solver. Thus, it is relatively easy to justify the local opportunistic maintenance search that occurs at the conclusion of Engineering Preventative Maintenance as a valid consequence of the physical system being diagnosed.

However, the search pre-constraining that occurs in Multinational Tax Planning (and the feedback that takes place to compensate for the shortcomings of such search) is very evidently a consequence of the inability of the human problem solver to know or simultaneously consider the whole domain of international tax regulations. Therefore, since KADS is targeted on the creation of a computer artefact, rather than a simulation of the expert, there is no reason why it should incorporate knowledge solely required to overcome the shortcomings of the human problem solver. Such components seem best left behind when the knowledge is transferred to the knowledge-based system and replaced by more rational structures.

Therefore, it would seem that KADS should not consider most of the problem solving components summarised in Chapter 3. Certainly we would not advocate the building of knowledge-based systems which merely mimicked the human problem solver's behaviour. (However, if this were required to be done then it would not appear to be difficult to construct interpretation models incorporating the problem solving components such as those discussed in Chapter 5.)

Nevertheless, although there is probably no merit in simulating such components, we feel that an understanding of their existence is rather useful. Such assertion is made on the basis of the findings on the worthwhile sample of problem solvers that we have investigated.

There appear to be three relevant conclusions. First, it seems likely that very similar components are found in experts solving, or attempting to solve, very different problems. Second, the problem solving of the experts is very much distorted, or non-optimised, to accommodate the components needed to overcome the expert's cognitive limitations. Third, some problems (such as Eurobond Warrant Trading) are very poorly "solved" because there is no real way that the problem can be addressed to make it solvable by the human "expert". (For instance, in the case of Eurobond Warrant Trading, the sub-domains are so inextricably linked that the human problem solver's strategy of serial sub-domain search is sensibly confounded.)

So, without overstressing the point, it could be useful for the KADS user to be at least aware that he or she may encounter knowledge in the human problem solver which specifically reflects the cognitive limitations of the human. Since components of such knowledge seems to be replicated in quite a small sample of problem solvers, it could be useful for KADS workers to be "warned" as to the sort of knowledge they might find and the ways in which they might recognise its presence. Perhaps recommendations could be made for interpretation models on which to "rationally reconstruct" the knowledge acquired from a non-optimal human problem solver. (Although an interpretation model was not used, this was exactly what happened in Multinational Tax Planning where the "rationally reconstructed" system as implemented did not constrain the problem to, say, South America but could take due consideration of all the companies in the world-wide group. This

383

lead to more comprehensive solutions without the need to feedback the constrained result for global checking.)

Thus prior awareness of problem solving components of the type highlighted in Chapter 5 could lessen the potential confusion of the KADS user and might lead more directly to the selection of better problem solving interpretation models. Finally, in those cases where the "expert" is actually not a very good problem solver, it might be possible to recognise this at an early state and, taking whatever knowledge the "expert" has, incorporate this in a knowledge-based system having none of the inherent limitations of the human problem solver. Thus, for instance, in Permanent Health Insurance, it was evident that the human expert just could not cope with all the permutations of Age, Health, Job, Leisure etc and, even if he could, he would have been unable to remember all the individual cases and combinations. Such limitation would not apply to a knowledge-based system to support Permanent Health Insurance and the KADS user, if he were to recognise a PHI or similar situation in advance, might be fairly confident that he could deliver a knowledge-based system that would outperform the human expert.

Finally, it is necessary to comment on the uptake of KADS in the expert systems community. KADS appears as a very useful and widely applicable methodology that has benefitted from the very adequate funding of a talented team of researchers and developers. So why has the author and, with very few exceptions, most of his commercial expert systems colleagues and competitors chosen not to use KADS for their everyday work? And why is it that toolkits like Shelley and KEATS are not in widespread use by commercial knowledge engineers?

The answer is certainly not that there is widespread criticism of KADS. Few expert systems practitioners have used or studied KADS and rejected in on grounds of principle. In fact, few expert systems practitioners have studied KADS at all. And of those that have made the attempt, even fewer (including the author) are fully confident that they understand it in detail. Until recently, the KADS enthusiast has had to undertake either a training course or disentangle a mass of research findings or ESPRIT reports. It is not straightforward to adopt KADS, even if there is willingness on the part of the knowledge engineer.

And, to be frank, most knowledge engineers seem unwilling. Many expert systems, although large, are quite straightforward to implement and knowledge engineers do not seem to need the elaborate KADS methodology to do their job. Even some experienced KADS practitioners will admit that they use little more than the Model of Expertise and that finding a ready-made appropriate interpretation model is a minority occurrence.

So KADS and the toolkits which support it are very much under-used in respect of the enormous intellectual investment made in them. This is a very unfortunate situation because KADS is well-founded and could satisfy the demands for the analysis, design and implementation rigour that is possibly preventing the uptake of knowledge-based systems in the more conventional software engineering community.

It is not the author's purpose to discover the means by which KADS (and Shelley and KEATS) might find wider usage. However, he will make a few further suggestions in addition to the ideas expressed above where he suggests that KADS ought at least to be aware of some of the more common peculiarities of the human problem solver.

385

First, KADS must be made more understandable and penetrable by the knowledge engineer. Otherwise, he will not use it and will carry on coping just like the experts that he studies. From the authors commercial experience he is aware that today's knowledge engineers are a very much more able and qualified than the run-of-the-mill software engineers who currently implement conventional software systems. If KADS is proving difficult for the dedicated expert systems practitioners, it will prove impossible for the systems analyst or programmer who has to be re-assigned to expert systems work. And, unless KADS or something similar can be used on the shopfloor there will be no widespread future for knowledge-based systems.

Even if mainstream developers were able to resort to, say, straightforward expert systems prototyping or incremental development (like many current knowledge engineers) IT management are now so convinced of the overriding need for software development methodologies and tools that these practices would be outlawed. So, if KADS is too difficult to use, then knowledge-engineering will not enter the mainstream. (There are actually further problems such as the reconciliation of KADS with object-oriented design and these issues are briefly discussed in the next section.)

Perhaps a mini-KADS is needed to enable simple tasks to be investigated and implemented with a minimum of knowledge of the full KADS methodology. Or, perhaps, KADS should be offered in task-specific versions. Diagnostic-KADS, Planning-KADS etc. These might offer an easier entry for the knowledge or software engineer with a requirement to service a specific commercial problem solving domain. And, offer a route for Chandrasekaran's (1987) Generic Task approach to find a proper and complementary role to the ideas of Wielinga and Breuker (1986) upon which KADS is based.

There may be a case for supporting the novice KADS practitioner with a knowledge-based system derived from the expertise from an experienced KADS practitioner. We remember the telling quotation from Wielinga, Schreiber & Breuker (1992) that knowledge engineers have to cope with a complex knowledge engineering process and that the availability of multiple models in KADS is an asset in this respect. Perhaps a further asset dedicated to the same end would be an inbuilt knowledge based system.

In summary, although the author did not use KADS in support of this research, it may yet remain one of the few routes to regularise the building of full or partial knowledge-based systems in the conventional software engineering community. There may be some utility for the KADS user to be able to recognise some of the problem solving components summarised in Chapter 5. However, any philosophical objections to KADS disappear in the author's main concern that, for the developer of average ability, it is just too difficult to use and understand.

7.4   KNOWLEDGE AS A COMPONENT OF CONVENTIONAL SYSTEMS - A NEW KADS ROLE?

"Conventional" and "Knowledge-Based" Analysis and Design

The following is a simplistic, but useful, comparison of conventional and knowledge-based systems analysis and design methods.

Conventional analysis and design methods increasingly look at the business model and strategy and so determine what computer functionality is needed

to support the business processes. Once the functional requirement is established then a hardware and software architecture is formulated that will satisfy that requirement.

Knowledge-based analysis and design methods, such as KADS, look at existing experts, their environments and interactions and attempt to model these. Once a model is established then a hardware and software architecture is formulated that will satisfy the model.

A Combined Approach

So, in outline terms, the conventional approach is fed from a business requirement and the KBS approach is based on the availability and exploitation of existing pockets of human knowledge. These approaches are different but not incompatible. In fact, at a fairly simplistic level, they combine very well.

A combined conventional/KBS approach would commence with the business model and its requirement for computer functionality to support it. In parallel the existing knowledge in the business (ie human experts) would be surveyed to determine the areas where there were knowledge-rich resources available.

Where the perceived knowledge would satisfy part of the required business-driven functionality then such knowledge would be used to support such functionality. Where knowledge was available that would satisfy functionality over and above the business-driven requirement then consideration would be made to increase such functionality if it further

satisfied the needs of the business.

In essence the combined approach would be based on a conventional analysis and design method but with an "opportunistic" element which would reveal, analyze, transform and incorporate available knowledge already employed in the business.

A Changed and Reduced Role for KADS

Such a combined approach would require some modification if KADS were to be used. In a normal KADS-supported KBS project an interpretation model is identified and the available knowledge is used to populate this model. Usually the model drives a top-down search for knowledge to be acquired from the expert(s). In a business-driven combined approach a model of required functionality would be identified and this would drive a top-down search for available knowledge to be used to populate the model. In such cases knowledge becomes a fortuitously available resource to answer some of the needs of the business-driven design.

If such a combined approach were to be formalised and taken beyond the very simple outline above then the role of KADS would require to be changed. Whereas the full KADS methodology would be used for "pure" KBS applications, it is probable that little more than the Model of Expertise would be required for the combined approach. For instance, the roles satisfied by the Organizational Model and the Model of Cooperation would not be required because, for instance, Context Diagrams, User Catalogues and Logical Dataflow Models would satisfy all or part of these roles. (It is to be remembered that it is really only the Model of Expertise that uniquely

distinguishes KADS from other methods. Most of the other KADS models have some sort of equivalent in conventional methods and, in a combined approach, such KADS models would be supplanted by their equivalents from the dominant conventional method.)

An Opportunity to Recover Lost Territory

However, although KADS would only be partially used, the combined approach could radically open up the opportunities for exploiting knowledge resources. The author agrees with Porter (1993) in that many of the most successful KBSs are those that are integrated with conventional systems rather than being stand-alone KBSs. However, and currently, the integration of a KBS with a conventional system is awkward at best and results in a heterogenous system, often with a discrete KBS "module" being embedded in and callable by a major conventional application. In the main, this does not attract users to KADS. In the main such conventional users would prefer and, because of their experience, find it easier to implement "knowledge" in conventional languages.

(The author does not want to enter a debate about what is knowledge or a knowledge-based system. However, where expert systems or KBS is not used, then substantial elements of what knowledge engineers call knowledge get incorporated, perhaps not very well, in conventional systems. The term "KBS" is always taken to refer to systems incorporating knowledge encoded in "AI" or "expert systems" languages. However, acquired knowledge can often, and often more efficiently, be encoded in conventional languages. A major UK bank has recently encoded a "knowledge-based" (ie based on acquired human knowledge) fraud detection system in the conventional

language PL/1. Is it a "knowledge-based system"? The author's opinion is that conventional business software has long incorporated knowledge to greater or lesser degrees. "KBS" or "expert systems" are systems where this is done "properly" using the methods and tools of the expert systems community.

The Closer Integration of Knowledge in "Conventional" Software

With the availability to survey and assimilate knowledge into an all-embracing conventional design we see the real possibility of the wider and more useful exploitation of knowledge in systems design and implementation. Such systems might not be entirely homogenous in their implementation in that, a rule language might be used to represent the knowledge and, for instance, an object-oriented language would encode the "conventional" parts of the system. However, since knowledge would be incorporated at the "fragment" rather then the "module" level, the overall integration of the knowledge would be much tighter than in those current systems where a "bolt-on" KBS module is employed.

Perhaps even more important, although knowledge would be acquired using expert systems (probably KADS) knowledge acquisition methods, the bulk of the systems design and implementation would be done using methods and tools (suitably enhanced where necessary) that were familiar to the conventional software community. A gradual transfer to systems explicitly incorporating knowledge could occur rather than the current situation where transfer to KBS technology is culturally difficult for the established software development community and, in the main, is ignored or does not occur.

As reported in Chapter 2, the KADS community are making KADS more robust, via CommonKADS (de Hoog, Martil, Wielinga, Taylor, Bright & van de Welde, 1992) so that it is more acceptable to those used to the rigour of conventional methods. However, KADS and CommonKADS remain primarily targeted at KBS developments. This section proposes methods and tools, based on conventional methods and tools but with additional coverage, perhaps KADS-provided, to cover the needs of developers wishing to undertake developments where knowledge is integrated at a low-level component level rather than a module level. The recent paper by Saward, Land & Bingham (1993) provides a good starting point for this process where it considers the requirements for a genuine integration of KADS with conventional software engineering methods.

7.5   KNOWLEDGE AS A REUSABLE DESIGNER COMPONENT

This short section is largely a rider to the previous section where knowledge is viewed in its role as a component of conventional systems. "Conventional systems" is, perhaps, a poor term since the world of conventional software is itself changing. The trend toward object-oriented systems is leading very much toward the encapsulation of useful functionality and its reuse in a myriad of different applications.

Perhaps this trend should put the genuine advocate of knowledge-based systems to shame. Because, in its origins as human knowledge, a primary characteristic of knowledge is its capacity for reuse. Remembering the domain of Benchtop Electronic Diagnosis, is it not the same knowledge about circuit behaviour that can allow an individual to, on separate occasions design, troubleshoot and extend the functionality of a specific circuit?

Knowledge of the characteristics of a particular chip are reused when deciding to employ that chip in the circuit or to hypothesise whether it is the cause of a particular fault.

It was the early dream of the advocates of production systems that knowledge, in the form of forward-chaining rules, would be floating ready in the system ready to progress any novel situation through to an unique conclusion. However, the reality is that expert systems and KBS developers have formulated their systems to reliably deal with quite tightly constrained types of problems. In particular, most expert systems replicate just a single type of task undertaken by an expert. Systems which can undertake, say, design, diagnosis and applications advice based on a common knowledgebase do not seem to exist. Yet a major percentage of the necessary knowledge is common to both requirements. (In the same way Life Underwriting and Permanent Health Insurance rely on much common knowledge, say, for interpreting the effect of smoking on an existing pulmonary health condition. Since an experienced human underwriter can address both tasks using elements from the same body of knowledge then it is regrettable that this cannot be done, and is not even considered, for expert systems to support the tasks.)

Given this background of poor reuse of knowledge, the integration of knowledge into conventional systems re-opens the opportunity to use knowledge acquired, say, by use of KADS as resource for incorporation in a number of related applications. Such applications would not be exclusively based on the particular tasks undertaken by the expert(s) from whom the knowledge was acquired. However, by identifying reusable (and reliable) packets of knowledge (perhaps encapsulated into or uniquely accessed by dedicated objects) the systems designer could use these as and when

required for current and future systems. (The word "packets" remains deliberately undefined but probably would correspond to the "sub-task" level in KADS. What is being proposed is essentially a library of reusable sub-task knowledge-based objects.)

Perhaps by taking the vehicle of reuse, this section has anticipated the role of acquired knowledge changing from a resource not solely dedicated to the construction of systems with similar overall functionality to those of the human experts from whom the knowledge was acquired. Packets of knowledge, once acquired and, very probably, tested for quality and reliability, could be used over and over again as part of the resource for the system designer.

7.6   SYSTEMS BEYOND THE LIMITATIONS OF HUMAN PROBLEM SOLVERS

In Chapter 3 we reported on investigations of a large series of problem solving tasks. As a consequence of those investigations we noticed a series of "problem solving components" which seemed to be related to the ways in which experts cope with complex or large problems. We separately discussed such components in Chapter 5.

One outcome of these considerations was that, although human experts manage to circumvent some of the limitations of their own cognitive abilities, there are certain tasks which are so particularly complex or large that such coping strategies are inadequate. Typically, although a task could be divided up into multiple sub-domains, the linking between the searches of such sub-domains was so heavy as to effectively frustrate the coping strategies that were used.

394

As a result, in such tasks as Permanent Health Insurance, Job Shop Scheduling and Eurobond Warrant Trading, the "expert" does not really produce optimum or reproducible solutions to the problems with which he is presented. Indeed, the term "expert" may be something of a courtesy title. Certainly the task holders have knowledge about their domain and can make reasonable inferences concerning elements of the problems which they face. However, they are often unable to adequately cope with the whole problem and this seems to disbar them from the title "expert".

Now, of course, there are other reasons why we could not be confident about supporting tasks like Eurobond Warrant Trading. However, there are other tasks which do not have these disbarring reasons (unstable knowledge, need to encode verbal and visual data etc) which could well be the subject of productive knowledge-based systems. In Chapter 6 it is suggested that, for instance, Permanent Health Insurance could be supported by KBS technology. However, it would require sponsor and expert to face the facts of the existing situation and contribute to a process where the corporate principles of underwriting such insurance were properly considered so that they could be incorporated into the design of comprehensive knowledge-based system.

Such a system, once constructed on the basis of explicit principles, would perform significantly better and more reliably than the human expert and would become an important asset to the company.

In effect our argument is that existing expert systems show no real advantage over human experts. In many cases, as discussed at the beginning of this chapter, expert systems cannot perform all the expert functions, or are costly to implement, and therefore sponsors will continue

to employ human practitioners as the most cost-effective option.

However, if expert systems (or more properly KBSs) can be constructed to perform well in those areas where human experts are at the edge of their powers, then their marketability will be greatly increased. For instance, although "heavy reference linking" may be a bad sign for the human problem solver, the computer is immune to such problems and, provided there are no other problems, an effective system can be constructed.

So, in summary, a key traditional rule for the intending expert systems builder was to base his system on acquirable knowledge from a competent expert. Therefore it may seem perverse to suggest that a good route to commercial KBS success might be to base systems on semi-acquired, semi-"designed" knowledge from an "expert" producing non-optimum solutions. However, providing the reasons for the non-optimality are understood (and here the considerations of Chapter 3 and 5 might help) and they are addressable by the technology, then there is a good chance of producing a commercially attractive and successful system.

7.7 SYSTEMS TO ADDRESS UNSATISFIED BUSINESS NEEDS

The last section has suggested that KBSs may be able to penetrate areas where, although there are alleged experts, such experts generally do not perform well.

This section, in contrast, illustrates such needs using a single example in the form of a very brief reminder of the possibilities for communicative

396

knowledge-based systems discussed in Chapter 6. Such systems do not pretend to replicate an existing problem solver, either good or poor, but suggest that existing human knowledge may be combined and supplemented to produce a novel communicative system. Of course this is just the issue discussed above under "Knowledge as a Reusable Designer Component". However, in this case we are adopting the sound conventional strategy of pointing to identifiable business needs (eg experts who do not talk to each other, identifying and using existing valuable knowledge is in the organisation) and proposing that these may be usefully and reliably satisfied by KBSs. Again, a commercially valuable functionality over and above that of any existing human problem solver would provide the developer not only with a justification for his work but also an opportunity not to be constrained by the lower costs of the competing human problem solver.

Although not discussed elsewhere in this thesis, recent commercial investigations by the author suggest that "corporate" knowledge-based systems, acting as accessible storehouses of both runnable knowledgebases and browsable textual knowledge, would provide a second and powerful example of knowledge-based systems answering the urgent and unsatisfied business need to encode, manage and make available valuable knowledge dispersed within a business or other organisation. Again such systems would provide a new and novel resource to an organisation that could not readily be provided by human experts.

7.8 SUMMARY - THE GENERAL PROSPECT FOR EXPERT SYSTEMS

The expert systems and KBS community enjoys many technically able and enthusiastic practitioners. The "products" are genuine and offer real

business advantages. Yet the "sales figures" do not justify the investment.

The author's belief is that we face a marketing failure. Expert systems have often been promoted for tasks for which they cannot succeed or are commercially uncompetitive. KBS methods and tools are increasingly effective but do not integrate well with conventional methods and tools. Yet the market requirements are for knowledge to be used as a component of larger systems. KBS methods and tools can often provide systems of similar competence to human experts but are not well suited to offering knowledge in smaller reusable packages for use by designers of new applications. Many systems have been produced which mimic existing cost-competitive and competent human experts. However, unique chances to address large and complex business problems which routinely confound human experts remain largely unexplored. The opportunities for the technology to enable business to realise and communicate its knowledge-based investments are hardly considered.

Fortunately marketing problems, once identified, can be subject to quite rapid resolution provided the underlying product technology is sound. The author is increasingly optimistic for the future of knowledge-based systems technology and hopes that the analyses in this chapter will contribute to the realisation of its long-awaited promise.

CHAPTER 8 - SUMMARY AND CONCLUSIONS

A Duality of Objectives

There has always been a duality in the objectives of this research. Although there are worthwhile studies of problem solving in, say, the managerial context (eg Wagner, 1991) or for "professionals" such as doctors (eg Elstein, Schulman & Sprafka, 1978) or lawyers (eg Amsel, Langer & Loutzenhiser, 1991) there are few studies of technicians or "semi-professionals" solving routine occupational problems. Similarly academic studies of problem solving seem so often to involve tasks which are either contrived or involve game playing or other worlds which are remote from the world of the workplace. Yet, for many workers, their ability to solve special sorts of problems is the attribute that decides whether they are to remain employable and what salary they may enjoy for their efforts.

So occupational problem solving is a serious activity and is a worthy objective for careful academic study. And, to be quite clear, such study does not require any connection or objectives to represent the discovered problem solving in any form of expert system. In a research context the sole use of any computer simulation would be to investigate and test hypotheses about the cognitive mechanisms underlying the discovered problem solving performance.

However, the other primary objective of the research is very much concerned with computer modelling of occupational problem solving, not for the purposes of cognitive simulation, but to enable expert systems to be built for commercial usage and profit. The research objectives for this

work were also very clear. How can knowledge be most efficiently acquired and implemented into a technically and commercially viable expert system? And, since this is the author's professional field, there has been an opportunity to follow both objectives in parallel. And, in doing so, get access to a range of problem solvers well beyond that available to the average academic researcher.

A Wide Range of Task Investigations

Problem solving for a wide range of tasks has been investigated. From, at one extreme, the workshop technician repairing electronic equipment to, at the other extreme, the world-class expert at optimising multinational taxation. But, because the research and investigation opportunities were subject to the uncertainties of the commercial situation, expert time was often limited and it was not always possible to carry through either investigations or implementations to a conclusion.

Constraints for the Author and the Experts

However, such commercial pressures provided the basis of many of the research findings. How could occupational problem solvers be investigated under real-life working conditions? How would methods like KADS have helped the author if he had had the courage and time to use them? What were the constraints that the problem solvers themselves suffered? The answers have ranged from lack of information in Benchtop Electronic Diagnosis, through to lack of management appreciation in Engineering Preventative Maintenance, through to "impossible" problems in Eurobond

400

Warrant Trading.

But, unlike their laboratory counterparts, occupational problem solvers cannot fail to deliver solutions. In Permanent Health Insurance this meant that the solution was really rather unprincipled. In Job Shop Scheduling it was just non-optimum.

The Components of Coping

And, in examining how problem solvers came to cope with difficult problems in the context of their own limited abilities, a set of "components" came to be recognised, occurring again and again in different tasks. Most were based on variants of a "divide-and-conquer" strategy but with a range of pre-and post-problem solving variations to constrain the problem or ensure that nothing had been missed as a result of such constraint.

A Miscellany of Reasons for Success or Failure

Further, the task investigations threw up a whole range of reasons why expert system might be feasible, or a success or a failure. And, in analyzing these reasons it became apparent why so many expert systems and KBS projects had not succeeded through to commercial success. Some were chasing knowledge that did not exist or was unstable. Some were addressing problems that nobody cared about. And some investigations were costing far more money than could ever be recovered in the "working life" of the expert system.

Finding Better Problems to Solve

Even the incomplete investigations were useful. In most cases a reason for the brevity of the investigation was that the sponsor did not think that there was a task to be addressed by the technology. But indeed there was. Several tasks were discovered where an organisation was struggling to understand and make available the knowledge and information which it possessed in its myriad of departments. And, in other cases it was obvious that problem solvers in separate departments could best communicate through knowledge-based systems proper to their mutual domain.

Serious Conclusions

And the real conclusions to this work were at the end of Chapter 7. And, as conclusions, they are very serious indeed. The UK government, the EC and industry has invested heavily in expert systems. Academics and businessmen have uniquely worked together to try to make expert systems a commercial success and an advantageous enabling technology for business. And, although it has not all failed, neither has it robustly succeeded.

Technology v. Accountancy

Perhaps it is a problem of an enterprise driven from the laboratory rather than the marketing department. There is no lack of good people or technology and most of us know that adequate techniques and toolkits have been with us for many years now and there is really no barrier to implementing an expert system where knowledge is available and adequate.

The problem is squarely in the understanding of the market.

Just because something is technically feasible, or even "clever" or "interesting", there is no reason why the hard-headed businessman (especially if he is an accountant) will buy it. Of course the government will put up pump-priming funds and industry will sometimes pay its share. But, unless expert systems can deliver real cost-effective business benefit then they will be politely and properly declined by industry.

So we can succeed with our able technology providing we take care to address it only to problems and opportunities where it has both clear technical and commercial advantage. This will cut down the number of candidate applications in the short term. But, as the applications filter through, the feedback to the market will bring forward the demand and opportunity that has so long been anticipated for the technology.

An Interim Report

Our journey is complete and there is much to do. This research has barely touched on the variety of occupational problem solving that warrants investigation. With a view to business, we have done little more than hint at the correct way to profitably deliver KBS technology to the market. So our original objectives remain with us and are unchanged.

Summary and Conclusions

ACKNOWLEDGEMENTS

This research has been aided by the advice, encouragement and co-operation of many individuals.

My colleagues in the Human Cognition Research Laboratory have provided me with a strong, friendly and enduring anchor into academic research. I particularly thank my patient supervisor Marc Eisenstadt and Hank Kahney and his family for their kindness, encouragement and advice.

British Telecom gave me the foundation and confidence to commence this work. David Memory's manipulation of bureaucratic procedures both provided me with a fine insight into the nature of occupational expertise and also the necessary authorization to embark on this research. George Pollard, my able friend and colleague on the great TRACKER adventure, was the first of the real computer wizards with whom I have been fortunate to work.

Two experts at British Telecom, Jim Harvey and Allan Ellis were the first unwitting victims of my research into knowledge acquisition. Both were genuine experts although unrecognised as such except by their immediate colleagues. It was the revelation of the extent and quality of their knowledge that convinced me of the potential value of expert systems in the everyday work situation.

Dick Stewart's robust faith in the future of AI was the key force that led Coopers & Lybrand into allowing me to undertake so much exciting work on their behalf. In the final audit his investment will be justified. Dieter

BIBLIOGRAPHY

Note: A small number of titles are included below which, although not specifically referenced in the text, have provided valuable background to the research.

Amsel E, Langer R and Loutzenhiser L (1991) Do lawyers reason differently from psychologists? A comparative design for studying expertise. In Sternberg R J and Frensch P A (Eds), Complex Problem Solving: Principles and Mechanics. Lawrence Erlbaum, Hillsdale, New Jersey.

Anderson J R (1976) Language, memory and thought. Lawrence Erlbaum, Hillsdale, New Jersey.

Anderson J R and Bower G H (1973) Human associative memory. Winston, Washington, D.C.

Anjewierden A, Wielemaker J et al (1992) Shelley – computer-aided knowledge engineering. Knowledge Acquisition, 4, 109–125.

Barr A and Feigenbaum E A (Eds.) (1981) The Handbook of Artificial Intelligence, Vol 1. Morgan Kaufmann, Los Altos CA.

Bennet J (1983) ROGET: a knowledge-based system for acquiring conceptual structure of a diagnostic expert system. Journal of Automated Reasoning, 1, 49–74

Bobrow D G, Kaplan R M, Kay M, Norman D A, Thompson H and Winograd T (1977) GUS, a frame-driven dialog system. Artificial Intelligence, 8, 155-173.

Bouchet C, Brunet E and Anjewierden A (1989) Shelley: An Integrated Workbench for KBS Development. In Proceedings of Expert Systems and Their Applications, Avignon.

Bramer M (1990) Practical Experience in Building Expert Systems. Wiley.

Breuker J and Wielinga B (1984) Techniques for Knowledge Elicitation and Analysis. ESPRIT Project P12, Report 1.5, University of Amsterdam.

Brown J S and Van Lehn K (1980) Repair Theory: A generative theory of bugs in procedural skills. Cognitive Science, 4, 379-426.

Bricklin D (1986) Demo Program User Manual. Software Garden., West Newton, MA.

Bundy A and Welham B (1981) Using Meta-level Inference for Selective Application of Multiple Rewrite Rule Sets in Algebraic Manipulation. Artificial Intelligence, 16, 189-212.

Burton A M and Shadbolt N R (1987) Knowledge Engineering. Technical Report 87-2-1, Artificial Intelligence Group, Department of Psychology, University of Nottingham.

Cendrowska J and Bramer M A (1982) A Rational Reconstruction of the MYCIN Consultation System. Open University

Chandrasekaran B (1987) Towards a Functional Architecture for Intelligence Based on Generic Information Processing Tasks. Proceedings 10th IJCAI, Milan, 1183-1192

Chase W G & Simon H A (1973) Perception in chess. Cognitive Psychology, 4, 55-81.

Chi M T H, Feltovitch P J and Glaser R (1980) Representations of physics knowledge by experts and novices. Report No.2, Learning Research and Development Center, University of Pittsburgh.

Clancey W J (1983) The Epistemology of a Rule-Based Expert System – a Framework for Explanation. Artificial Intelligence, 20, 215-251.

Clancey W J (1985) Heuristic Classification. Artificial Intelligence, 27, 289-350.

Carroll J M & McKendree J (1987)Interface design issues for advice-giving expert systems. Communications of the ACM, 30, 14-31.

Davis R (1976) Applications of meta level knowledge to the construction, maintenance and use of large knowledge bases. Report No STAN-CS-76-564, Stanford University.

Davis R (1979) Interactive Transfer of Expertise: Acquisition of New Inference Rules. Artificial Intelligence, 12, 121-157.

Davis R (1980a) Meta-Rules: Reasoning about Control. Artificial Intelligence, 15, 179-222.

Davis R (1980b) Content Reference, Reasoning about Rules. Artificial Intelligence, 15, 223–239.

Davis R, Buchanan B and Shortliffe E (1977) Production Rules as a Representation for a Knowledge-Based Consultation Program. Artificial Intelligence, 8, 15–45.

de Hoog R, Martil R, Wielinga B, Taylor R, Bright C and van de Velde W (1992) The CommonKADS method. ESPRIT Project P5248 KADS-II, Technical Report KADS-II/WPI-II/RR/UvA/018/4.0.

Dreyfus L and Dreyfus E D (1986) Mind Over Machine. Blackwell, Oxford

Dym C L (1987) Issues in the design and implementation of expert systems. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 1(1), 37–46.

Elstein A S, Schulman L S and Sprafka S A (1978) Medical Problem Solving: An Analysis of Clinical Reasoning. Harvard, Cambridge, Mass

Ericsson K A and Simon H A (1980) Verbal Reports as Data. Psychological Review, 87, 3.

Erman L D and Lesser V R (1975) A multi-level organisation for problem solving using many, diverse, co-operating sources of knowledge. IJCAI-4, 483–490

Eshelman L, Ehret D, McDermott J and Tan M (1988) MOLE: a tenacious knowledge acquisition tool. In Boose J H and Gaines B R, Knowledge

Based Systems, 2, 95-108, Academic Press, London.

Falk G (1970) Computer interpretation of imperfect line data. AI Memo 132, Stanford University.

Feigenbaum E A and McCorduck P (1984) The Fifth Generation, Pan Books, London.

Gaines B and Linster M (1990) Development of Second Generation Knowledge Acquisition Systems. In Wielinga B, Boose J, Gaines B, Schreiber G and van Someren M, eds., Current Trends in Knowledge Acquisition, 143-60, Amsterdam.

Gammack J G and Young R M (1985) Psychological techniques for eliciting expert knowledge. In Bramer M A (Ed.) Research and Development in Expert Systems, 105-112, Cambridge University Press, Cambridge.

Gammack J G (1987) Different techniques and different aspects on declarative knowledge. In Kidd A L (Ed.) Knowledge acquisition for expert systems: a practical handbook, 137-164, Plenum, New York, NY.

Georgeff M P (1982) Procedural Control in Production Systems. Artificial Intelligence, 18, 175-201.

Gleeson J F J and West M L J (1988) CLINTE: Coopers & Lybrand International Tax Expert System. In Moralee D S (Ed.) Research and development in expert systems IV, 18-31, Cambridge University Press, Cambridge.

Hayes-Roth F, Waterman D A and Lenat, D (1983) Building Expert Systems. Addison-Wesley, Reading MA.

Hershey D A, Walsh D A, Read S J and Chulef A-S (1990) The effects of expertise on financial problem solving: Evidence for goal-directed, problem-solving scripts. Organizational Behaviour and Human Decision Processes, 46(1), 77-101.

Hesketh P and Barrett T (1989) M1 An Introduction to the KADS Methodology. Esprit Project P1098 Report, STC Technology Ltd, Harlow, Essex.

Hewitt C (1972) Description and theoretical analysis (using schemata) of PLANNER, a language for proving theorems and manipulating models in a robot. Report No TR285, AI Laboratory, MIT.

Jackson P (1990) Introduction to Expert Systems. Addison-Wesley, Wokingham.

Kahney H and Eisenstadt M (1982) Programmer's mental models of their programming tasks. Proceedings of the Fourth Annual Meeting of the Cognitive Science Society, 143-145.

Kidd A L (Ed.) (1987) Knowledge Acquisition for Expert Systems: A Practical Handbook, Plenum, New York, NY.

Laublet P, Mepham M, Motta E, Zackland M, Eisenstadt M, Gerders M, Rugg G and O'Hara K (1992) Review of Software Mediated Knowledge Acquisition. VITAL Internal Deliverable ONERA/T211.4/I/1.

412

Lenat D B (1982) The Nature of Heuristics. Artificial Intelligence, 19, 189–249.

McDermott D (1978) Planning and Acting. Cognitive Science 2, 71–109

Michalski R S (1974) Variable-valued logic: system VL1. In Proceedings of Fourth International Symposium on Multiple Valued Logic, Morgantown, West Virginia.

Michalski R S and Chilausky R L (1980) Knowledge Acquisition by encoding expert rules versus computer induction from examples: a case study involving soyabean pathology, Int J Man-Machine Studies, 12, 63–87.

Michalski R S and Larson J B (1978) Selection of most representative training examples and incremental generation of VL1 hypotheses: The underlying methodology and description of programs ESEL and AQ11. Report 867, Computer Science Department, University of Illinois.

Miller G A (1956) The magical number seven, plus or minus two: Some limits of our capacity for processing information. Psychological Review 63, 81–97

Minsky M (1975) A Framework for Representing Knowledge. In The Psychology of Computer Vision, Winston P H, McGraw-Hill, N Y, N Y.

Mittal S and Dym C L (1985) Knowledge acquisition from multiple experts. AI Magazine, 5(2), 32–36.

Motta E, Rajan T and Eisenstadt M (1988) A methodology and tool for

knowledge acquisition in KEATS-2. Technical Report 32, Open University.

Musen M A, Fagan L M, Combs D M and Shortliffe E H (1988) Use of a domain model to drive an interactive knowledge editing tool. In Boose J H and Gaines B R, Knowledge Based Systems, 2, 257–273. Academic Press, London.

Neale I M (1988) First generation expert systems: a review of knowledge acquisition methodologies. Knowledge Engineering Review, 3(2), 105–45.

Nelson T H (1965) A file structure for the complex, the changing and the indeterminate. Proceedings of the ACM National Conference, 84–100.

Nelson T H (1973) A Conceptual Framework for Man–Machine Everything. Proceedings of the National Computer Conference, M21–M26.

Newell A (1973) Production systems: Models of control structure. In Chase W (Ed), Visual Information Processing, Academic Press, New York, 463–526.

Newell A (1982) The Knowledge Level. Artificial Intelligence, 18, 87–127.

Newell A & Simon H A (1972) Human Problem Solving. Prentice Hall, NJ.

Nilsson N J (1971) Problem Solving Methods in Artificial Intelligence. McGraw-Hill, New York, NY.

Nisbett R E and Wilson T D (1977) Telling more than we can know: Verbal

reports on mental processes. Psychological Review, 84, 231-259.

Norman (1983) Design Principles for Human-Computer Interfaces. In Baeker M & Buxton A S (1983) Readings in Human-Computer Interaction: A Multidisciplinary Approach. Kaufmann, Los Altos, California.

Norman D A, Rumelhart D E and the LNR Research Group (1975) Explorations in Cognition. Freeman, San Francisco.

Porter D (1993) KADS-II: Towards the CommonKADS method. ESPRIT Project P5248 KADS-II, Technical Report KADS-II/T5.1,2/PP/TRMC/001/1.0.

Pople H E Jr. (1982) Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnosis. In Szolovits P (Ed.) (1982) Artificial Intelligence in Medicine. Westview Press, Boulder CO.

Post E (1943) Formal reductions of the general combinatorial problem. American Journal of Mathematics, 65, 197-268.

Rieger C (1976) An Organization of Knowledge for Problem Solving and Language Comprehension. Artificial Intelligence, 7, 89-127.

Russo J E, Johnson E J and Stephens D L (1989) The validity of verbal protocols. Memory and Cognition 17(6), 759-769.

Saward G, Land L and Bingham T (1993) Framework for Method Integration (Product Level). ESPRIT Project P5248 KADS-II, Technical Report KADS-II/TM9.1/TRMC/011/0.4.

Schank R C and Abelson R P (1977) Scripts, plans, goals, and understanding. Lawrence Erlbaum, Hillsdale, NJ.

Schreiber G (Ed.) (1992) Special Issue: The KADS approach to knowledge engineering. Knowledge Acquisition, 4 (1).

Shortliffe E H (1976) Computer-based medical consultations: MYCIN. American Elsevier, New York.

Shaw M L (1982) Attending to Multiple Sources of Information: 1. The Integration of Information in Decision Making. Cognitive Psychology, 14, 353-409.

Smith S L & Mosier J N (1986) Guidelines for Designing User Interface Software. Bedford, MA: Mitre Corporation. (Report #MTR-10090.)

Soloway E, Ehrlich K, Bonar J and Greenspan J (1982) What do novices know about programming? In Badre A and Shneiderman B (Eds.) Directions in human-computer interactions, 27-54, Ablex, Norwood, NJ.

Steele, G.L (1984) Common Lisp: The Language Digital Press, Burlington, MA

Steels L (1988) Components of Expertise. AI Memo 88-16, AI Lab, Vrije Universiteit Brussel.

Szolovits P and Pauker S G (1978) Categorical and Probabilistic Reasoning in Medical Diagnosis. Artificial Intelligence, 11, 115-144.

van Melle W, Shortliffe E H and Buchanan B G (1981) EMYCIN: A

Domain-Independent System that Aids in Constructing Knowledge-Based Consultation Programs. Infotech State of the Art Report.

Wagner R K (1991) Managerial Problem Solving, In Sternberg R J and Frensch P A (Eds), Complex Problem Solving: Principles and Mechanics, Lawrence Erlbaum, Hillsdale, New Jersey

Wason P C and Johnson-Laird P N (1972) Psychology of Reasoning: Structure and content. Harvard University Press, Cambridge, Mass.

Waterman D A and Hayes-Roth F (1978) Pattern-directed inference systems, 483-501 Academic Press, New York.

Weiner J L (1980) BLAH, A system which explains its reasoning. Artificial Intelligence, 15, 19-48.

Weiss S M, Kulikowski C A, Amarel S and Safir A (1978) A model-based method for computer-aided medical decision-making. Artificial Intelligence, 11, 145-172.

Welbank M (1983) A review of knowledge acquisition techniques for expert systems. Martlesham Consultancy Services, Ipswich.

Wielinga B J and Breuker J (1986) Models of Expertise. Proceedings ECAI'86, 306-318.

Wielinga B J, Schreiber A Th and Breuker J A (1992) KADS: a modelling approach to knowledge engineering. Knowledge Acquisition 4(1), 5-54.

Wielinga B J and Shadbolt N (1989) Conceptualisation of a Knowledge Engineering Workbench. Deliverable 1.4, Esprit Project P2576, University of Amsterdam.

Williges B H & Williges RC (1984) Dialogue design considerations for interactive computer systems. In Muckler F A (Ed.), Human Factors Review, Human Factors Society, Santa Monica, CA.

Winston P H (1978) Learning by Creating Transfer Frames. Artificial Intelligence, 10, 147–172.

APPENDIX A – TELECOM REPAIR SERVICE: ANALYSIS OF ACQUIRED KNOWLEDGE

A.1   TELECOM REPAIR SERVICE RAW PROTOCOL TRANSCRIPT

(The extract below is taken from lines 1–110 of a transcript of 1296 lines.)

BARNET REPAIR SERVICE CONTROL – FIRST SESSION (EXTRACT)

E = Chas Church

S = Jim Harvey

1   E   So when there's not much information on the docket you either go back or what?

2   S   Yes, normally you do your basic testing and make sure everything looks OK, then you've got to ring the customer and see if you can get the information you want from them. It comes again more difficult because if you get no reply then, from our point of view here you just mark it 'Testing OK; no reply' and it'll go down the other end; a chap down the other end will try it again later on perhaps several times during the day to try and get an answer to see if we can get any more information.

3   E   This is one that worries me a bit, where you've got a customer report and it's testing OK what sort of faults

could still exist that you can't see?

4 S   You can get a situation where you pick the phone up without getting a condition on the phone, we don't get a loop from the tone, you can't tell that unless they can answer you or try to answer you, or unless you know someone's there ...

5 E   In other words you can see the capacitor?

6 S   You can see the capacitor but when they pick it up you've still got the capacitor; you see it an happen. A simple answer is the plungers are stuck down. So you can get a situation there where they say I've got no dialling tone and if they've only got one instrument –

        if you say to them is there someone there and they say yes and you ring on it and you can't get any answer then you've got to accept that there is something of that nature wrong with that end: it may not be the plungers stuck, I mean it could be something else.

7 E   But normally in those cases what would you do? If you got just an 000 fault report.

8 S   If we get that sort of report and it's come especially from the manual board our instructions are to re-test it and if after 24 hours we can't get an answer from the customer we write it off as RWT NRC which is no reply from customer.

9   E   But really that excludes all those bits of faults ...

10   S   That's right, yes. We do get a few but we run a system where if a customer comes back within a couple of days and says I'm still having trouble and you look at the card and you say oh yes, he reported it and we RWT NRC, then it will go to the front of the pile so that he doesn't suffer too much.

11   E   That's really the only thing you can do under those circumstances apart from send a chap out.

12   S   That's right. Well they found that we were having so many abortive visits because the number that came back as against the number that were written off was very small.

13   E   What sort of things can you see at the customer apparatus end without the customer actually co-operating in testing? Is there anything you can see at all?

14   S   The sort of things we can see, generally, are faults that could be not just in the customer's premises unfortunately, it could be on the line itself outside. You can get a rectified loop or ...

15   E   Where would a rectified loop come from?

16   S   You can get them in the premises or externally, or in odd

occasions in the exchange – it's unusual in the exchange but you can get them there.

17 E   How would you rectify that?

18 S   It's usually caused through dampness: if you check on here,instead of getting with an ordinary R resistance loop in that situation where you're feeding out a battery on the B–wire and we've got an earth on the A–wire so if there's any connection between the A and B we'll get a deflection. With an ordinary loop condition – I mean I can show you if you like, that's probably the easiest thing to do. You see in that condition we've got a deflection and we know that it is a loop.

19 E   So you're now saying you've got what: sort of 30K on that,is that right? On the ohm scale?

20 S   Yes, and if I reverse the line so that the battery goes out on the A–wire and the earth is on the B the needle stays in the same place but obviously with a rectified loop that will either disappear completely or it will go a long way back.

21 E   So it's not full rectified, it's a change?

22 S   It's a change, yes, but it's a term ... I mean a lot of people years ago when I first came on the job when I was on external I was always taught initially that it was

called a low loop; now if you've ever heard of a low loop, it was just a term used, but if somebody said to you it was a low loop you'd think ah, yes, I've got to be careful here because when you're using a manual tester - it's all right sitting here you just automatically flick the keys but when I was outside we used to have a Detector 4 which you just had a couple of wires that you could play around with. One way you get a capacity but if you turn them round you'll get the rectified loop come up.

23 E   The Tester SA9083 does that - we'll talk about that sometime.

24 S   Because you've got keys on there and you can do it but where you've got small testers ... so it's helpful to the people outside if you say to them it's a rectified loop because they know that it's no good just sticking a meter on it and looking at it one way round, you've got to reverse it.

25 E   So when it's rectified in the network it's what I call half=and-half; when it's in the sub's premises is it full rectified, I mean plonk and nothing?

26 S   No, no, it varies. We might come across some actually if we do some. I'll see whether I can find any. That's going to be a battery I think, that one but I'll try it anyway.

27 E   Why do you think that's going to be a battery?

28 S   It just happens that I know that we've got a breakdown
               in that area so the chances are ... but it's possibly a
               mains one.

29 E   So you've started your diagnosis even before you've
               done the tests?

30 S   That's right, yes. The customer rung me up and he said
               he's getting problems coming in ...

31 E   You know where it is by the postal address or what?

32 S   Yes, he said to me he's having trouble, I said what road
               are you in, he said Warren Road, I said we've got a
               problem down there – I know that we've got a major
               breakdown developing in that area. Yes, it's going to be
               a battery and there it is.

33 E   So he's got a battery on what: on the B leg?

34 S   Yes, he's got a 50-volt battery on the B-wire and on the
               A-wire also we've got a battery but it's dropping off
               which indicates that it's a damp one.

35 E   That battery on the A-leg is that coming via the loop or
               is it coming via where?

36   S   No, it's just picking it up. What's happened is the cable is wet so it's picking it up off other circuits.

37   E   Yes but what I'm saying is the B-leg's got full battery more-or-less, now the A-leg is that picking it off independently from somewhere else not via the ...

38   S   Probably, probably, I can't tell.

39   E   Because there's no loop there, so it couldn't pick it through the loop could it?

40   S   It could do but it's difficult to tell in this instance because there's very little on the loop scope because that really is the difference between my meter and the battery that's on there. Really what we're looking at the moment is just the deflection: I know it's a battery but I was just looking to see whether ... well you can actually see what sort of condition you get when there is a rectified loop on the line.

41   E   So we're looking at the loop are we now?

42   S   So you've got 19 divisions on that one, then that way round we've got 35, you know, it's altering; but I know it's not a rectified loop because it's battery A and B, but that's the sort of thing ...

43   E   I get what you mean, because there's volts on it you're

not really testing a loop.

44 S   This is where you've got to be careful because you've got to be positive as to what you're looking at. I mean when you throw that key and you look at it to me that's just a deflection, it means nothing at the moment but I've got to throw that key to decide, now, see, it's a battery – it's on the move the whole time. Probably if I listen on it it's as noisy as anything. Yes.

45 E   That's water you reckon or what?

46 S   That is, yes.

47 E   So listening's a pretty good test for that is it?

48 S   You can usually tell, I mean I would never bother to prove that out in an exchange because in normal circumstances if you get a full battery like that, in the majority of cases if you're getting a full 50 volts on a line if you apply an earth to it and listen usually, if it's going to be in the exchange, you'll get dialling tone as a general rule. I mean unless you've got an absolute full 50 volts – OK sometimes the exchange battery to us reads 48 volts, as long as you know what the exchange voltage is then unless it reads 48, if it's less than 48 or 50 or whatever you're almost 100% certain if you say it's away from there.

426

49 E So anything below about 45 you know it's out of the exchange.

50 S OK, you're going to be wrong by probably one in perhaps several hundred.

51 E How could it be? You've got a leaky path in the exchange?

52 S It's possible but the chances are very very remote: it's like I said I mean I have known a rectified loop being in the exchange but when you find it eventually you find it was in a jumper that had got warm with other jumpers running across it and it just sort of arced across: it's just one of those things, so really when you're doing this there's a sort of calculated risk in other words.

53 E Do you get much feedback? I mean that one you said about a jumper, how would you eventually know that?

54 S You've given it out to an external chap and he'd say well I'm proving it back to the exchange and you say OK I'll get the fuses pulled, you get the fuses pulled and you can see it. Usually what I do is get the chap in the exchange to put this where he's in on the other side of the main frame and if it's in that jumper then you'll get a clear user dis, but if it's bad to there it's going to be up on the IDF so it's just then dealt with by the special faults chap.

55   E   I mustn't hold you up Jim you carry on and we'll just go through a routine with you.   So that was what you thought it was,probably in that cable.

56   S   Yes, that's right, that's where it's going to be: there's no doubt about that one. This one's a bit different but ...

57   E   What have we got there?

58   S   She says that on her first line that she's not getting calls in. I did ask her if she could get dialling tone on it but she said I don't know.

59   E   That's a PBX fault is it?

60   S   Yes. It's blocked up in the exchange, it's held up by a previous incoming call.

61   E   How does the test tell you that?

62   S   We've got PET condition: well if we press the prior control key down we then override the PET condition and we're getting 50 volts, I know that's going to be the exchange battery. If we plug into another one, 49 volts, it's the same, and if we check the other leg,there's nothing on there except the exchange earth and there's no dialling tone on there, and if I disconnect my loop which I've got on the line at the moment you can just

sort of hear the relay click in and out.

63 E   Which relay's that?

64 S   That'll be one in the exchange; probably in the final. It's going to be that this PET in the exchange is held up.

65 E   Where would that PET come from in the exchange?

66 S   It's possible that it could be something which is not supposed to happen nowadays and that is that what we call the CSH, it's Called Sub Held, it shouldn't happen officially.

67 E   Because it should be sub clear down shouldn't it?

68 S   That's right: it should be automatic clear down on next meter pulse; but I mean I took that call at 18 minutes past nine and here we are at 5 past 11, two hours later and it's still there.

69 E   So it could be that?

70 S   It could be or it's possible that there's a fault at the distant end which is holding the call

71 E   In other words the calling subscriber is still really holding it.

72 S   That's right. It might not be coming right from the other phone but it's being held, I would say, probably from the calling exchange.

73 E   But you could hear the exchange?

74 S   Yes. This one's a little bit different.

75 E   A single business line is it?

76 S   No, it's a Herald this one. She complains that calls coming in are cutting off. Let's put a call in over the network and see what happens.

77 E   You operate his private control do you?

78 S   No, I've released it from my equipment altogether and I'm just putting a normal call out to them, I want to see whether ... just to see. Good morning, it's the telecom engineer. You complained a bit earlier on about incoming calls cutting off; could you tell me does this happen directly you've answered or is it after you've been speaking for some time? Is this after you've put it through to an extension or as soon as you answer, or ...? No, you shouldn't be able to lose a call, that's true, not from your end. But you've had it happen where you've extended them to an extension and then it's dropped back to you at the board has it? It's not come back to

you at all, they've lost it completely. All right, well we'll have a look: I don't think it's going to be anything at your end though, by the sound of it; whether we've got a problem somewhere else, I don't know. We'll make some tests, thankyou, bye. I don't like that one: that's not a nice one. They're getting cut off on incoming calls: well, basically, they can't really lose a call,especially on a Herald.

79 E   She's getting cut off on all sorts of conditions is she?

80 S   Yes, she says she gets cut off sometimes when she's talking to them before she's extended them to an extension and then when she extends to an extension she suddenly sees the light come up and the call's lost and she can see that the extension is flashing trying to get re-connected but the call itself is lost. If it gets cut off from the extension it should, on a Herald, drop back to the board but it's not happening so she's losing it completely which means really that there's... What it could be - and this is one of the problems of this - what it more than likely is probably a see first in another exchange somewhere which is dropping out on metering; now it could be on the initial meter pulse or it could be on subsequent meter pulses.

81 E   But wouldn't that depend on her getting the calls, all these ones from the same place?

431

82 S   That's right. It's very unlikely that they're all coming from the same place which again makes you wonder exactly where it is. It's a difficult one to find that one, I'm afraid, but I think I'll have to give that one to the man down on the section. I can't really give him any other information other than that really; I'll speak to him when I godown there about it.

83 E   On the section you mean ...?

84 S   It'll go down to the DFDR.

85 E   So you've more or less cleared that's exchange.

86 S   I would say so but the chances of finding it at the moment are rather remote unless we can pin-point that they all are coming from one particular exchange which is difficult because she doesn't know which calls are going to cut off; so unless you get her to say where are you calling from on every one ...

87 E   Which you can't do can you?

88 S   No, that's what I mean. This is what makes that sort of thing ... the diagnosis of that very difficult. I mean you know virtually that it's not at their end ...

89 E   So really you've got a fair idea of what the fault might be but you don't know how to find where it is – two

432

different things aren't they?

90 S   That's right. I mean I could be wrong. It could even be something silly like somebody's changed the wires on the final selector and reversed them, but that would cut her off on answering so it's not going to be that because she said it happens also when she's had it extended.

91 E   And you reckon the first meter pulse is the one that ...

92 S   Probably, subsequently. This is the sort of thing that we get.

93 E   What have we got here?

94 S   Shared service line: the wire sub complains of no dial tone. Again, we've got a PET condition. I'll hazard a guess that one of them has gone TOS.

95 E   Can you see both batteries or what?

96 S   I can see both of them now actually.

97 E   Was that testing normally?

98 S   Yes. I think we'll have to consult the TOS book on that one.

99 E   Would the other sub TOS affect it?

100 S   It shouldn't do but if it's done wrong it does.

101 E   You mean at the exchange they've plugged up the whole ...?

102 S   That's right, you've got both of them out instead of one. I may be wrong but that to me is the first thing to check.

103 E   It means if the other sub hasn't paid their bill, they cut both of by mistake.

104 S   One of the problems being with things like that is that when the exchange clerk calls over lines to be TOS they don't always tell you that it's a shared service line, so the chap in the exchange just goes round, moves the link and you've got both of them out.

This ones just a coinbox complaining that they get continuous pay tone after they've put 5p in. All we can do is just test the line just to see whether there is any fault. The sort of thing that can cause that on the line is,again, if you've got a rectified loop or a battery on the line.

105 E   Because it can't see the proper signal because that's a high resistance loop isn't it?

106 S   It's testing perfectly OK that one. We'll have to send out

on that one it could be a fault on the box or it could be in the exchange.

107 E   On the fee checker, yes.   But I imagine it's at the coinbox you notice it.

108 S   Yes, we get more problems, especially, as I say, 5p doesn't work but the 10p does, though the chances of it being on the coin feecheck in the exchange is a bit remote.

109 E   That wouldn't extend would it?

110 S   No, that's right.

END OF TRANSCRIPT EXTRACT

A2   TELECOM REPAIR SERVICE PROTOCOL DIGEST

(This digest corresponds to the extract from the raw protocol at Section A.1)

TRANSCRIPT DIGEST FOR BARNET REPAIR SERVICE CONTROL

JIM HARVEY – FIRST SESSION – RCO DUTY

1	If not much on docket then basic testing plus ring customer.
2	If customer no reply then mark card testing OK – no reply and try later.
3	What faults cannot see if testing OK but customer (valid) report?
4	Must have customer there to know if they can establish loop when called.
5-6	Still see capacitor when customer picks up phone.
	If NDT + one instrument + customer there + no reply then plungers stuck or something else.
7-8	OOO (especially from manual board) + re-test after 24 hours and no reply then write off RWT NRC.

9-10   If RWT NRC and one of few that come back then top of
            pile.

11-2   Only alternative to RWT NRC is send faultsman usually
            abortive. Number of RWT NRCs that come back is very
            small.

13   What can be seen at customers premises without
            customer's cooperation?

14   Unfortunately usually can only see things which could be
            either in customer's premises or on outside line.

14-6   Rectified loops from premises or externally or unusually
            from exchange.

18   Rectified loops usually caused by dampness.

18-20   Ordinary loop with battery on B-wire and earth on A-
            wire gives 30k deflection but reversing line on rectified
            loop (battery on A, earth on B will cause deflection to go
            right back or disappear.

22   Old term for rectified loop is low loop.

22   On external tester must reverse wires to see rectified
            loop with capacity one way and rectified loop the other.

24	Telling outside faultsman it is a rectified loop reminds them to reverse meter.
25-26	Rectified loops at premises or network are not different.
26-32	Battery fault assumed before testing on basis of customer difficulty "coming in" and postal address near major breakdown.
28	Major breakdown faults may be mains ones.
32	Expected battery found on "major breakdown" fault.
32	50-volt battery on B-wire, battery dropping off on A-wire indicates damp one.
35-40	Battery on A-leg is picked off other circuits and probably no loop there.
41-2	Testing loop "resistance" gives 19 and reverse 35 divisions not rectified loop because battery A and B.
42-44	Deflection altering..."on the move the whole time" will be battery so assume and then confirm "as noisy as anything".
44	A loop test is "just a deflection" and doesn't mean anything so must be positive what you are looking at. Confirm by battery test.

438

45-7   Confirmed as water - usually tell by listening.

48-50   Normally if 50 volts (and generally dialling tone) on line when apply earth to it then assume in exchange. If anything less than 50 volts (except where exchange battery reads 48 volts) then almost 100% certain away from exchange.

50-52   Could be wrong on less than 50 volts in exchange in one of several hundred faults by (rectified) loop in "arcing" warm jumper with other jumpers running across it.

54   Cases such as rectified loop in exchange noticed when given to external man and he asks for fuses to be pulled to prove back to exchange (after failing to find external fault?). If loop still seen when fuses out then dis wedges on other (exchange) side of MDF. If loop clears then (ie clear user dis) then in MDF jumper field.If bad to dis wedges then fault on IDF side and dealt with by Special Faults.

55-6   Fault at (32) confirmed as in cable "no doubt about that one".

58-9   PBX report of no incoming on first line. Customer not sure if getting DT.

60   Diagnosed as "blocked up in exchange - held up by

previous call".

61-4
Diagnosed by getting PET on line and then operating Private Control Key to override PET gives 50 volts on B-leg (confirmed as exchange battery by comparison (=49 volts) with another line and exchange earth and no dialling tone on A-leg. Disconnecting test loop can just hear (probably exchange final) relay click in and out.

65-73
PET suggested as due CSH which should clear down on next meter pulse but actually two hours since report. Suggested as fault at distant exchange although probably not distant customer.

74-6
Report of Herald with incoming calls cutting off.

76-78
Releases test desk circuits and puts in normal call over network "to see what happens".

78
Speaks to customer and asks if happens directly call answered or after speaking for some time, if when put through to extension or as soon as operator answers. Asks operator if had it happened where call has been extended to extension and the dropped back to board. Operator confirms not dropped back but lost completely. Tells operator fault not likely to be at her end and promises tests.

79-80
Confirms that operator is losing calls when speaking

before extending and when call extended (extension flashes for re-connection). If Herald gets cut off at extension then should drop back to board but this case is where losing call completely. Suggested as C1st on another exchange dropping out on metering (either initial or subsequent meter pulses).

81-9   Suggested diagnosis would depend on all calls from same exchange which makes "wonder exactly what it is". Difficult to find so given to FDO and will speak to him about it. Chances of finding fault are remote since customer does not know which calls will cut off and from where. However virtually certain not at customers end. Case where can diagnose probable cause of fault but not location.

90   Rejected alternative hypothesis is somebody changing and reversing wires on final selector which would cause cut off on answering. Actual fault also occurs when extended so fault on first meter pulse probable.

128   (Comment on Herald report)

93-4   NDT report from Y-sub on SS line. PET condition found. Guess that one sub has gone TOS.

95-8   Both batteries can be seen and line is testing normally (for SS).

441

98/102   Decides to (later) consult TOS book. May be wrong but first thing to check.

99-104   Other sub TOS should have no effect but exchange clerk may call over lines to be TOS without mentioning SS so exchange just move link and both subs out.

104-6   Coinbox report of CPT after 5p inserted. Only option is to test line which proves OK so send faultsman out. Could be fault in box or exchange.

104   Fault of this type could have been caused (and detected) by rectified loop or battery on line.

107-10   More problems at coinbox than at exchange fee checking equipment. Unlikely to be fee checker since fault is only on 5p and not 10p.

END OF TRANSCRIPT DIGEST EXTRACT

A.3   TELECOM REPAIR SERVICE KNOWLEDGE ORGANISATION – INTRODUCTION

This section includes extracts from the various knowledge organisation texts and structures used for analysis of the Telecom Repair Service knowledge. The extracts are largely self-explanatory and commence after the construction of the protocol digest at Section A.2 above.   Section A.4 illustrates the extraction of key words from the transcript digest and Section A.5 shows how the various domain concepts are ordered into a higher level classification.   Section A.6 is a hierarchical classification of concepts.

Section A.7 provides an example of concepts analysis and leads to Section A.8 which shows an actual example of analysed concepts.   At Section A.9 an extract is made of diagnostic sequences and consequential rule sequences are shown at Section A.10.   This final appendix includes examples from the initial rule set which formed the basis for the actual implementation.

A.4   KEY WORDS FROM TRANSCRIPT/DIGEST

(Page numbers refer to the original full transcript)

PAGE 1

docket

basic testing

ring customer

customer = sub

CNR = customer no reply

mark card

testing OK

RWT = right when tested

try later

report

valid report

customer there

establish loop

loop

see capacitor

capacitor

pick up phone (ie customer)

NDT = no dial tone

one instrument

(two instruments)

instrument = tele = telephone

tele

telephone

plungers stuck

plungers

OOO

manual board

re-test after 24 hours

top of pile

send faultsman = visit customer

visit customer

abortive visit

coming back (eg RWT NRCs)

seeing (conditions on test)

customer's premises = premises

premises

outside line

rectified loop

loop

external

exchange

cause

dampness

ordinary loop

battery B

earth A

put battery on B

put earth on A

reversing line

put battery on A

put earth on B

deflection

30k deflection

low loop = rectified loop

external tester

reversing wires (in field)

one way

the other way

capacity = capacitance

reversing meter (in field)

network

battery fault

testing

customer difficulty

customer I/C difficulty

postal address

major breakdown

mains

mains faults

expecting battery

PAGE 2

50 volt battery on B wire

battery dropping off on A wire

damp fault

battery on A wire

leg = wire

battery picked off other circuits

446

no loop

loop resistance

reverse loop resistance

battery A and B

deflection on the move

noise

noisy

loop test

battery test

water (=damp=dampness?)

listening

DT = dialling tone

apply earth = put earth

50 volts

48 volts

away from exchange

in exchange

arcing in jumpers

jumper

jumper fault

rectified loop in exchange

external man

pulling fuses

proving back to exchange

external fault

fuses

fuses out

dis wedges

(exchange) side of MDF

loop clearing

MDF jumper fault

MDF jumper

IDF side of MDF (=exchange side of MDF)

Special Faults

fault in cable

PBX

first line

no incoming

DT = dial tone

CSH = called sub held

blocked up in exchange = CSH

PET = permanent engaged tone

private control key

overriding PET

50 volts on B leg

measuring exchange battery

50 volts exchange battery testing another line

exchange earth

NDT on A leg

disconnecting test loop

hearing relay click in and out

CSH clear on next meter pulse

(elapsed) time

fault at distant exchange

Herald

incoming calls cutting off

normal call over network

speaking to customer

448

asking what happens

(fault) directly on answering

(fault) after elapsed time

putting through to extension

PAGE 3

(PBX) operator answering incoming call

(PBX) call extended to extension

(PBX) call dropping back to operator

(PBX) call lost completely

promising tests (to customer)

(PBX) calls before extending

(PBX) calls after extending

(PBX) board = operator

extension flashing

exchange C1st dropping out on metering

calls from same (distant) exchange

FDO

speaking to FDO

chances of finding fault

remote chance of finding fault

customer knowing source of calls

customer not knowing source of calls

fault at customers end

cause of fault

location of fault

reversal of wires on final selector (fault)

cut off on answering

fault on first meter pulse

Y sub on SS line

X sub on SS line

TOS = temporary out of service

seeing both batteries on SS line

SS line testing normally

TOS book

checking/consulting TOS book

first thing to check

effect of TOS on other sub

SS by moving link on IDF

TOS fault due to moving SS link on IDF

CPT after insertion of coin

fault in coinbox

fault in fee-checking equipment

fault on accepting any coin

fault on accepting particular coin(s)

1,000 ohm (RSC) loop

350 ohm cord (RSC) loop

650 ohm (RSC) loop (old exchanges)

using 1,000 ohm loop

using 650 ohm loop

using 350 ohm cord loop

END OF EXTRACT

A.5   CLASSIFICATION OF CONCEPTS

The following is a high-level classification of concepts. The labels in capital letters (eg FINDINGS) refer to separate concept listings for which an example (ELECTFIND) is given later in this appendix.

CONCEPTS IN RSC DIAGNOSIS

TESTS: deliberate information-seeking procedures (eg testing another line)

FINDINGS: resulting directly or indirectly from test procedures

REPORTS: reports from the customer or via operator (eg NDT)

CUSTQEST: diagnostic questions asked of the customer (eg incoming calls only)

CUSTTELL: customer information resulting directly or indirectly from questions (eg bedroom tele is on hook)

SEERECS: information seeking from paper or computer records (eg in TOS book?)

RECSINFO: records information (eg three reports I/C diff in last three months)

DIAGS: the diagnosis of the fault source (eg rectified loop in

premises)

DISTRIB: where or who the fault is distributed to (eg SFs, faultsman jointer)

FAULT: the eventual precise location of the fault (eg MDF fuse mounting, dial)

CLEAR: the "coded" record of the fault for future reference (eg dis in cabinet)

STRATEGIES: diagnostic strategies (eg prove out of exchange)

STRATEST: diagnostic strategies during testing (eg asking for fuses out)

ELECTEST: electrical tests done from the Test Desk (eg loop test)

ELECTFIND: results of electrical tests done from the Test Desk (eg battery B)

LISNTEST: listening tests done from the Test Desk (eg listening for birds)

LISNFIND: results of listening tests done from the Test Desk (eg modem tone)

TONETEST: listening for tones (eg SS dial tone by applying earth)

TONEFIND: results of tone tests (eg PET)

NOISTEST: listening for non-tone "noises" (eg listening to faultsman's short)

NOISFIND: results of "noise" tests (eg relay clicking in and out)

TONEFIND:

A.6   CONCEPT HIERARCHIES IN TELECOM REPAIR SERVICE DIAGNOSIS

TESTS

- ELECTEST

- LISNTEST

- TONETEST

- NOISTEST

- CUSTQEST

- SEERECS

FINDINGS

- ELECTFIND

- TONEFIND

- NOISFIND

- CUSTELL

- RECSINFO

REPORT-> [STRATEGY-> TESTS-> FINDINGS-> DIAGNOSIS]-> CLEAR

A.7   EXAMPLE OF CONCEPTS ANALYSIS

The following is an extract from one of the concept listings (ELECTFIND) shown in Section A.6 above.   Page numbers refer to the original full transcript

ELECTFIND

PAGE 1

testing OK

loop

see capacitor

NDT = no dial tone...NOISRESL?

seeing (conditions on test)

rectified loop

loop

ordinary loop

battery B

earth A

deflection

30k deflection

low loop = rectified loop

capacity = capacitance

battery fault

mains

PAGE 2

50 volt battery on B wire

battery dropping off on A wire

damp fault

battery on A wire

no loop

loop resistance.....ELECTEST

reverse loop resistance....ELECTEST

battery A and B

deflection on the move

50 volts

48 volts

loop clearing

50 volts on B leg

50 volts exchange battery

exchange earth

NDT on A leg

PAGE 3

cut off on answering

seeing both batteries on SS line

SS line testing normally

1,000 ohm (RSC) loop

350 ohm cord (RSC) loop

650 ohm (RSC) loop (old exchanges)

456

A.8   ANALYSED CONCEPTS

The following is a re-analysis of concepts from the REPORTS and DIAGS (DIAGNOSES) categories.

INITIAL CLASSIFICATION OF REPORTS

General and DEL Reports

- OOO

O/G O/G difficulty

- NDT
- NDT after making call
- locked up and not releasing after O/G call (= PET after O/G)
- reversion to DT
- CSH

I/C I/C difficulty

- no incoming
- incoming calls cutting off
- cut off on answering (= ring trip)
- outgoing but not incoming
- hearing ring current on answering
- intermittent bell ringing

I/C PET

- PET after O/G call

457

- ring trip

- CNI problems

Xmsn

- bad transmission

- reception fading

- faint reception

- interrupted DT during conversation

Noise

- Noise

- Noisy transmission

- hum

- O/G hum

Xtalk

- crossed lines

- conversation

Bell

- OOO

- bell tinkle

- faint

Misc

- new line fault

- non-complaints with massive earths and noise

- reports with lines testing perfect

Reports particular to general customer apparatus

- dial stuck

- "tele jobs"

Reports particular to SS

- NDT on Xsub

- disputed usage

Reports particular to PBXs

- first line

- Herald

- K&LU

- Monarch self-diagnostic reports

- extension fault

Reports particular to PCOs

- CPT

- CPT after insertion of coin (=CPT?)

- reversion to DT

- forced release

- COAMI (inaccurate PCO report)

- slots jammed

- foreign coins (= report?)

- door not closing

Reports particular to PWs

Operator Reports

- NU no record

- TNR (Test No Reply)

- (CNI)

Customer Service Requests

- putting NU on long standing fault

Engineering Requests

- LIAR reports

INITIAL CLASSIFICATION OF DIAGNOSES

Testing

- OK

- RWT

- RWT when tested with sub

- RWT because nothing can be done

- RWT when TNR but wife at home

- TOS

- PFT (locking up and not releasing) after O/G due bad battery A

- PET on SS due IDF link

- loop

Loop

- testing loop (eg receiver off)

- loop in customer's premises

- ordinary loop (= loop)

- no loop

- resistance loop (ie abnormal loop resistance) rectified loop

- in exchange

- in UG

- in premises

- rectified loop leading to ring trip low loop (= rectified loop)

- capacity

Dis

- intermittent dis
- dis A
- dis A (= 5 divisions)
- dis in exchange
- dis B at exchange
- dis outside
- dis between cabinet and BT (for jointer)
- dis end SS (cannot test)
- HR dis
- HR fault on line

Battery

- battery fault
- battery noise
- resistance battery
- leaky battery
- low battery fault in damp older cabinet assembly low battery fault rejected by SALT but within limits
- metallic fault = full 50 volt fault full contact with another line = 50 volt fault battery B
  - battery B giving ring trip
  - 50 volts battery B
  - battery B on PCO causing COAMI
  - battery A
  - nasty battery A
  - bad battery A causing PET/not releasing after O/Less serious

battery A leading to noise battery A and B
  - local battery

462

- positive battery

- positive battery from private equipment

- positive battery from telex

Earth

- earth fault

- earth A

- 23 ohm earth A

- earth A preventing ring trip

- earth A allowing O/G calls but not I/C

- earth A causing hum on O/G calls

- low resistance earth A from damp pole-top DP HR earth A

- earth B giving PET and NDT

- earth A and B

Ringing

- ring trip due rectified loop

Mains

Noise

- noisy transmission

- battery noise

- noise due to less serious battery A

- noise from cabinet

- noise from joint

- noise due HR joint

- noisy transmitter

Hum

- earth hum

- earth A hum when speaking to customer

- mains hum

Tele

- plungers stuck

- dial stuck

- dial fault (speed and digits)

- receiver off fault

- tele fault not requiring line test etc

Attachments

- answering machine ring trip

- answering machine full ring trip and tone modem ring trip and tone

- call diverter

Cause

Damp

- damp fault

- dampness

- water

Full contact

- exchange

- cabinet

- joint

- line

<u>Site</u>

- network

- testing centre spurious ringing

- distant exchange

    - routiner ringing fault

    - C1st dropping out on metering

    - crossed lines

- exchange

    - stuck N1st misrouting fault

    - crossed lines

    - loss of exchange DT

    - coin and fee checking equipment

    - group selector out

    - multiple PBX lines on one group error

    - uniselector fault

    - L&K relay fault

    - CNI

    - moving link on IDF for SS TOS (causing PET)

    - reversal of wires on final selector

    - MDF exchange side (IDF side)

    - jumper fault

    - cutting down to provide new line error

    - cutting down without applying NU error

    - full contact fault

    - dis in exchange

    - rectified loop in exchange

    - MDF line side

- fuse mounting

- fuses

- old-type circular fuse displacement

- fuses out = no fuse in

- discarded jumper wire

- solder droppings

- dis wedges

- dis wedges after meter tests

- dis wedges for CNI

- dis B at exchange

- dis on MDF or fuse

External

- major breakdown

- fault in 400 pair

- fault in 300 pair

- cabinet

- major cabinet fault

- cabinet assembly fault

- low battery fault in damp older cabinet assembly

- intermittent full contact at cabinet

- dis at cabinet

- intermittent dis at cabinet

- dis between cabinet and sub's joint

- cable breakdown fault

- fault in cable

- fault in 50 pair

- rectified loop in UG

- dis in joint

- HR joint

- noise due HR joint

- (PBX ringing lead continuous ringing fault)

- (pillar fault)

- outside line

- end SS dis (inability to test)

- DP fault

- DP fault with existing same DP fault

- low resistance earth A from damp pole-top DP customer's premises

- premises (= customer's premises)

- fault at customer's end (= premises)

- internal wiring

- rectified loop

- dis bell circuit

- multiple teles on wrong jacks

- telephone

        - receiver off

        - bell physical fault

        - transmitter fault

        - noise

PBX etc

- HES exchange jacks on windowsill

- local battery faults

- HES internal battery fault from other B wires

- battery fault from PBX switchboard

- K&LU odd battery and earth conditions

- Monarch self-diagnostic conditions

- extension fault

<u>Call Office</u>

- coinbox fault

- slots jammed

- foreign coins fault

- battery B causing COAMI

- door not closing customer

A.9   DIAGNOSTIC SEQUENCES

The following illustrates the way in which sequences of actions are extracted prior to the isolation of rules

TYPICAL RSC DIAGNOSTIC SEQUENCES

General beginning sequence for many fault conditions:-

Read docket -> Scan Card -> Electrical Tests
        -> Aural Tests -> Ring Customer etc

BASIC LINE TEST

Throw monitor...........Initial listen on line   KZ momentary

Throw voltmeter key......Tests B leg for earth   KVM

Throw receive key........Tests B leg for (-ve) battery   KVM and KRN

Throw line reverse key...Tests A leg for (-ve) battery   KVM, KRN and
                                                              KRA

Restore receive key......Tests A leg for earth   KVM and KRA

Throw earth key...........Earths A leg

Restore line reverse key..Tests capacitance   KVM, KE +/- KRA

BASIC LINE TEST FOR ENHANCEMENT

Throw monitor...........Initial listen on line   KZ momentary

(a) Throw voltmeter key......Tests B leg for earth   KVM

(f) Throw receive neg key....Tests B leg for (-ve) battery   KVM and KRN

    (Instructions suggest done after loop test)

(g) Throw line reverse key...Tests A leg for (-ve) battery   KVM, KRN and

    KRA

(b) Restore receive key.......Tests A leg for earth   KVM and KRA

(e/c) Throw earth key...........Earths A leg   KVM ,KE and KRA

    Tests (reverse) loop..drops back for rectified

(e/c) Restore line reverse key....Tests capacitance   KVM, KE +/- KRA

    Tests normal loop

If loop throw speak and battery key KKA not monitor KZ

A.10   OUTLINE RULE SEQUENCES

The   following   are   typical   rule   sequences,   assembled   on   cards,   and
supporting the discovered diagnostic sequences.

SEQUENCE 1

start   => read docket + scan card

NDT after O/G + card scanned   => test line

NDT after O/G + HR Earth A   => FDO/SF then C

SEQUENCE 2

start   => read docket + scan card

NDT + card scanned   => test line

NDT + TPET   => read exchange on card

NDT + PET + Exch 440   => check TOS book

TOS   => EFA

SEQUENCE 3

start   => read docket + scan card

NDT + card scanned   => test line

NDT + low C   => test C each leg

A and/or B leg low C   => FDO/dis out

SEQUENCE 4

start	=> read docket + scan card
PET + card scanned	=> test line
PET + loop	=> listen for noise or sounds
loop + sounds	=> apply howler
clear after howler	=> call customer + explain + EF

SEQUENCE 5

start	=> read docket + scan card
NDT + card scanned	=> test line
NDT + line tested	=> PC key
PC key + NDT + DT	=> 1000 ohm key
NDT + DT + 1000 ohm key	=> contact customer
NDT + customer DT restored	=> re-file as "customer misoperation"

SEQUENCE 6

start	=> read docket + scan card
NDT + card scanned	=> test line
NDT + low C	=> test C each leg
NDT + v.low C leg A + normal C leg B	=> test C alt junction
v.low C leg A - C alt junction < 2 divisions	=> FDO/dis B in exchange

472

SEQUENCE 7

start   => read docket + scan card

NDT + near bad cable   => test line

batt B + NDT + TNU   => UG fault + advise customer

APPENDIX B – SUMMARY CHARACTERISTICS OF TASKS INVESTIGATED

COMMERCIAL BUSINESS UNDERWRITING (Insurance Company)

Task Description	:Analyse business risk and prepare quotation
Implementation Details	:CUES using AION ADS+IMS+IBM 3090
Problem Descriptio	:Invitation to quote from branch
Prob Struct Aids	:Freehand A4 Basis of Quotation – Virtual Form
Generic Data	:Commercial Risk Manuals, Special Tables
Specific Data	:Previous quotes
KSs	:Fully linked PL,EL, Rating, Discounting
Open/Closed Problem	:'Closed'
Proactive/Reactive	:Reactive to Invitation to Quote
Problem Solving	:Simple Search Tree with Test at nodes and leaves. Wide and shallow search.
Comment	:Search may sometimes be constrained by business type coverage (eg PL, EL) specified by invitation to quote

FOREIGN EXCHANGE EXPOSURE (Aero Engine Manufacturer)

Task Description	:Devise method to avoid FX exposure
Implementation Details	:PC Crystal
Problem Description	:Telephone or written request
Prob Struct Aids	:Standard form of calculation
Generic Data	:Reuters, Telerate etc
Specific Data	:Request for exposure cover
KSs	:1
Open/Closed Problem	:Closed
Proactive/Reactive	:Reactive to Purchasing Department request
Problem Solving	:Simple Search (Algorithmic) Tree – Test at nodes and leaves
Comment	:Minimal expertise

BENCHTOP ELECTRONIC DIAGNOSIS (British Telecom)

Task Description	:Locate fault in piece of electronic equipment
Implementation Details	:PC Based TRACKER in micro-Prolog and μ-Lisp
Problem Description	:Fault docket from field faultsman
Prob Struct Aids	:Notes
Generic Data	:Notes (Diary), diagram, description, manuals
KSs	:(1) Search and Test
Open/Closed Problem	:Closed – finite number of components
Proactive/Reactive	:Reactive to "Faulty"/specific description
Problem Solving	:Dynamic Search Tree – Test at nodes/leaves Diagnostic identification at leaves Proceduralised hypothesis and test at nodes (optimised search path). Unexpected clues dynamically rearrange or prune tree. Initial report also rearranges or prunes extent of tree (eg bell fault).
Comment	:1 sub-domain fully searched at 'Faulty' default case

AUDIT WORK PLANNING (Coopers & Lybrand)

Outline Description	:Analyse audit requirement & make work plan
Implementation Details	:Expertest using QShell in GC Lisp
Problem Description	:Audit partner's report
Prob Struct Aids	:Audit schedule form
Generic Data	:Audit manual
Specific Data	:Audit File including previous years schedules
KSs	:Individual procedures + limited combination
Open/Closed Problem	:Closed
Proactive/Reactive	:Reactive to annual statutory need
Problem Solving	:Multiple Simple Search with limited overarching (conciliator) findings
Comment	:1 sub-domain fully searched

BUSINESS PERFORMANCE ASSESSMENT (Government Health Department)

Task Description	:Make a report on an operational health unit
Implementation Details	:DH Performance Advisor Crystal/Lotus 1-2-3
Problem Description	:None
Prob Struct Aids	:Draft -> Final Segmented Report
Generic Data	:PI Relationship Charts
Specific Data	:District Data + Objectives, Previous Reports
KSs	:(1)
Open/Closed Problem	:Closed
Proactive/Reactive	:Reactive to Departmental Need to Report
Problem Solving	:Multiple Simple Search (semi-algorithmic) with Secondary Grand Conclusion
Comment	:Essentially exhaustive search

MULTINATIONAL TAX PLANNING (Chemical Manufacturer)

Task Description	:Find a tax-efficient corporate structure
Implementation Details	:MNCT using ART/Joshua/Lisp
Problem Description	:Memo from Corporate Treasury
Prob Struct Aids	:Trial partial structures, trial schedules
Generic Data	:Existing structure + fiscal regulations
Specific Data	:Subsidiary balance sheets + schedules
KSs	:(2) Group viewpoint + company viewpoint
Open/Closed Problem	:Open
Proactive/Reactive	:Reactive to acquisition, dividend request etc Proactive for more tax-efficient structure
Problem Solving	:Pre-Constrained Search Tree (or forward chaining) with local result affecting overall result. Constraints consist of semi-permanent constraints on holding relationships plus 'artificial' problem definition constraints to current focus of attention company. Internalised multi-level feedback (via calculation) to check effect on overall tax paid on dividends to parent company. Can be considered as damped iterative search.
Comment	:Local results fedback for global consideration

ENGINEERING PREVENTATIVE MAINTENANCE (Cigarette Manufacturer)

Task Description	:Locate mechanical fault & make repair plan
Implementation Details	:Not implemented
Problem Description	:Verbal breakdown report
Prob Struct Aids	:Work schedules
Generic Data	:Maintenance schedules (specific assemblies)
Specific Data	:Machine maintenance record
KSs	:(2) Planned + Breakdown
Open/Closed Problem	:Closed
Proactive/Reactive	:Reactive breakdown
	:Proactive opportunistic maintenance
Problem Solving	:Pre-Constrained Search Tree (or forward chaining) without local result affecting overall result but localised opportunistic forward chaining. Initial fault report prunes tree but does not rearrange it.
Comment	:Search constrained to breakdown assembly – necessary repair access causes search for opportunistic maintenance

PERMANENT HEALTH INSURANCE (Insurance Company)

Task Description	:Analyse a health risk and prepare a quotation
Implementation Details	:Not implemented
Problem Description	:PHI Application Form
Prob Struct Aids	:Freehand A4 Basis of Quotation
Generic Data	:Library of previous quotes, Special Rates
Specific Data	:Application Forms, Medical Reports
KSs	:Health x Occupational Risk x Occupational Need x Money -> Combinatorial Evaluation
Open/Closed Problem	:Closed
Proactive/Reactive	:Reactive to Invitation to Quote
Problem Solving	:Multiple Wide Shallow Scan with Massive Combinatorial Grand Conclusion (ie Multiple Diagnostic Search with Forward Chaining Combinatorial Conclusion)
Comment	:Search as with Life but wider coverage

LIFE UNDERWRITING (Insurance Company)

Task Description	:Analyse a life risk and prepare a quotation
Implementation Details	:New Life ES, AION ADS, IMS, IBM 3090
Problem Description	:Life Application Form
Prob Struct Aids	:NOBI 'Form'
Generic Data	:Life Tables+Special Rates+Cultural stereotypes
Specific Data	:Application form, Medical reports
KSs	:Actuarial+Medical+Occupation/Residence/Sport
Open/Closed Problem	:"Closed" – artificially constrained
Proactive/Reactive	:Reactive to 'problem' on application form
Problem Solving	:Multiple Wide Shallow Scan (or forward chaining) with Institutionalised Constrained Conclusion
Comment	:Search constrained by application form data in three fixed sub-domains

SMALL BUSINESS GUIDANCE (Government Training Agency)

Task Description	:Make a report on a small business
Implementation Details	:TAP Small Business Advisor
Problem Description	:Company name & business activity description
Prob Struct Aids	:Interview notes
Generic Data	:Sector reports. Financial newspapers/journals
Specific Data	:Accounts + company reports
KSs	:Multiple (7)
Open/Closed Problem	:Open
Proactive/Reactive	:Proactive to better business goal.
Problem Solving	:Multiple with Minimal Reference Linking
Comment	:Initial diagnostic conversation/analysis

PENSIONS MANAGEMENT (Public Utility Pension Scheme)

Task Description	:Answer specific question from scheme member
Implementation Details	:PC Crystal
Problem Description	:Verbal or written enquiry from member
Prob Struct Aids	:Letter to member
Generic Data	:Pensions Manual, HQ Advisory Circulars
Specific Data	:Member's records and correspondence
KSs	:
Open/Closed Problem	:'Closed'
Proactive/Reactive	:Proactive to enquiry from member
Problem Solving	:Pre-Constrained Search with Minimal Inter-Access to Other Search Areas
Comment	:Implementation problems due to multiple conflicting experts

JOB SHOP SCHEDULING (Consumer Products Manufacturer)

Task Description	:Scheduling multi-stage processes in a job shop
Implementation Details	:Not implemented
Problem Description	:Order data etc
Prob Struct Aids	:Gantt charts, schedules, capacity plan
Generic Data	:Machine capacities, maintenance plans
Specific Data	:Order data, existing schedules
KSs	:Multiple
Open/Closed Problem	:Closed (but enormous solution space)
Proactive/Reactive	:Reactive to monthly scheduling requirement
Problem Solving	:Multiple Simple Search with Heavy Interlocking – Fixed Set of Sub-Domains
Comment	:Poor 'expert' performance

TELECOM REPAIR SERVICE (British Telecom)

Task Description	:Locate fault in a local telephone network
Implementation Details	:"DIS" Lisp-based ES
Problem Description	:Fault docket
Prob Struct Aids	:Docket
Generic Data	:Tables, Cable Fault Records
Specific Data	:Customer Fault History
KSs	:(3 constant) Report, Test, Fault History
Open/Closed Problem	:"Closed"
Proactive/Reactive	:Reactive to problem description
Problem Solving	:Cooperative Search with Heavy Linking
Comment	:Sometimes sub-domain constraint from docket

EUROBOND WARRANT TRADING (International Bank)

Task Description	:Fix a Eurobond option price
Implementation Details	:Not implemented
Problem Description	:None
Prob Struct Aids	:Slip, Spreadsheet
Generic Data	:Reuters (+News), Telerate, Financial Times etc
Specific Data	:Reuters (specific page)
KSs	:Multiple
Open/Closed Problem	:Open
Proactive/Reactive	:Reactive to Need to Buy/Sell or answer request for quote Proactive to 'Make Money' goal
Problem Solving	:Multiple Simple Search with Heavy Interlocking - Variable Set of Sub-Domains
Comment	:Reactive fairly successful; proactive gut feel - "52% successful"

480

APPENDIX C – SUMMARY OF TASKS AND PROBLEM SOLVING COMPONENTS

TASKS

Telecom Repair Service
Pensions Management
Eurobond Warrant Trading
Job Shop Scheduling
Small Business Guidance
Life Underwriting
Permanent Health Insurance
Engineering Preventative Maintenance
Multinational Tax Planning
Business Performance Assessment
Audit Work Planning
Benchtop Electronic Diagnosis
Foreign Exchange Exposure
Commercial Business Underwriting

PROBLEM SOLVING COMPONENTS

Global Data Acquisition

Initial Problem Analysis

Pre-Constraining

    – Institutionalised

    – Domain-Driven

Single Domain

Tests at Intermediate Nodes

Multiple Sub-Domains

    – Discrete

    – Overlapping

    – Fixed Set

    – Variable Set

Sub-Domain Reference Linking

    – None

    – Minimal

    – Heavy

    – Problem Solving

    – Hypothesis Strengthening